Heterogeneous Effects of Regulatory Enforcement: Evidence from OSHA

Taylor Watson *

November 12, 2025

Latest version available here

Abstract

Opponents of regulation argue it stifles economic activity — but what actually happens when governments enforce regulations on noncompliant firms? I estimate how enforcement affects individual plants in the context of the US Occupational Safety & Health Administration, linking data on surprise inspections to restricted-use Census micro data. Using a matched-stacked event study design, I find that while high regulatory fines decrease employment and increase closure probability, low fines have the opposite effect. I interpret these results with a conceptual framework in which plants try to comply with those regulations they believe to be profitable, while a regulator uses fines to incentivize compliance. The framework predicts that fines are lower when safety is more profitable, and higher when the social benefits exceed the private costs. I provide evidence consistent with this prediction, showing that high fines are associated with higher private costs such as reputation damage and productivity losses, while low fines may result in profitable spillovers such as improvements in management. These findings show that enforcement effects are not uniformly negative, and may in some cases improve plant performance.

*Boston University, watson@bu.edu

I am indebted to Kevin Lang, Matt Johnson, Tarek Hassan, Keith Ericson, and Johannes Schmieder for exceptional guidance, and to the participants of the Boston University Empirical Micro Seminar for valuable feedback. I am also grateful to Shital Sharma and Sari Kerr at the Census Boston Research Data Center for indefatigable assistance in data access. All errors are my own. This project uses data from the U.S. Census Bureau. Any views expressed are those of the authors and not those of the U.S. Census Bureau. The Census Bureau has reviewed this data product to ensure appropriate access, use, and disclosure avoidance protection of the confidential source data used to produce this product. This research was performed at a Federal Statistical Research Data Center under FSRDC Project Number 3216. (CBDRB-FY25-0446). This research uses data from the Census Bureau's Longitudinal Employer-Household Dynamics Program, which was partially supported by the following National Science Foundation Grants SES-9978093, SES-0339191 and ITR-0427889; National Institute on Aging Grant AG018854; and grants from the Alfred P. Sloan Foundation.

I Introduction

Opponents of business regulation argue it is burdensome, costing workers their jobs and reducing firm profits (Branstetter et al. 2013; Alfaro-Urena et al. 2022; US Chamber of Commerce 2025). Some firms choose not to comply with regulations, preferring to risk fines rather than bear such burdens (Stansbury 2021). If the government intervenes at these firms to enforce regulations, this argument would predict decreases in plant performance — yet little evidence exists on how regulatory enforcement affects noncompliant plants.

This paper provides new evidence on the effects of regulatory enforcement, using the context of the US Occupational Safety & Health Administration (OSHA). Among plants receiving surprise inspections from OSHA, I compare those whose inspections resulted in fines (enforcement) to those that did not. I link these plants to restricted-use Census microdata to estimate the longitudinal effects on employment, survival, and worker earnings.

I find that the effects of enforcement are not uniformly negative for plant performance. Plants whose violations result in low fines (below the average of \$100 per worker) see an increase in employment of around 5% relative to unfined plants. Similarly, low-fined plants experience a decrease in the probability of closure of roughly 0.5 percentage points (a 50% decline), suggesting a meaningful difference in profitability post-inspection. In contrast, plants receiving high fines (over \$100 per worker) see a roughly 1% decrease in employment, and a sharp 70% increase in closure probability.

To ensure causal identification, I propensity score match plants that receive fines with highly-similar unfined plants, then use a matched-stacked event study design to estimate effects. The resulting estimates are causal under the identification assumption that the conditional difference between fined and unfined plants would have been stable in the absence of fines. I defend this assumption through pre-trend analysis and contextual details from OSHA's inspection process, which minimizes the differences among inspected plants.

To understand these heterogeneous effects, I develop a conceptual framework that generalizes safety choices over heterogeneous workplace hazards. Profit-maximizing plants remediate hazards when they believe the benefit (avoided workers' compensation and potential fines) exceeds the costs (safer production processes). With imperfect information on these factors, they sometimes fail to make profitable remediations. Plants are then visited by a surprise regulator, who enforces remediation of all hazards. The effect of this enforcement depends on the efficiency of the hazards being remediated. If remediation is privately efficient, then enforcement can benefit the plant; if not, then enforcement has negative effects.

To illustrate the differences in efficiency across hazards, consider two chemicals commonly used in construction: lead and methylene chloride. Lead exposure causes serious diseases, but can be safely mitigated if clothing is properly laundered. In contrast, exposure to methylene chloride has serious, but quite rare, consequences, and is quite costly to mitigate — requiring expensive air filtration systems and monitoring. As a result, enforced remediation of lead has minimal profitability effects (and may actually improve it), while enforced methylene chloride remediation will involve high costs and minimal benefits to the plant (Occupational Safety and Health Administration 2010, 2025a).

These differences in private efficiency across hazards imply that the regulator can use different levels of fines to incentivize compliance. For those hazards which are less privately costly (or even profitable), the regulator can apply low fines and still make the expected value of compliance greater than that of risking a fine. For those hazards which are more costly — where the private efficiency is low but social benefits are high — the regulator must apply higher fines in order to ensure compliance. For the above example, the average fine for lead violations is $12 per worker, compared with $6,500 for methylene chloride.

This model prediction of different fines across different hazards is supported by the empirical heterogeneity analysis. Low fines are associated with improved performance, while high fines have negative effects. To further understand what drives these effects, I explore the differences between private and social effects of remediation, both in benefits and costs.

While the private and social benefits of safety are similar at low levels of fines, at higher levels the private benefits plateau while social benefits continue to rise. The primary benefit of remediation is avoided injuries, which are socially valuable to workers and privately valuable to the firm through avoided workers' compensation. The key factor that OSHA considers in setting fines is the potential risk facing workers. This creates a positive correlation between fines and both private and social benefits; however, the private benefits are often capped by statutory maximums on compensation payouts and legal frictions (Oliphant and Wagner 2012). This means that for sufficiently high fines, the social benefits far outweigh the private benefits. Additional private benefits exist in the form of spillovers, as argued in Haviland et al. (2010) and Berman and Bui (2001), including potentially through improved management (Bloom et al. 2025).

I argue that private costs are higher for those plants receiving high fines, through effects on productivity and in the labor market. Remediations associated with high fines are associated with productivity losses in related settings, as shown by Li (2022), suggesting similar effects for safety. I show that firms hit with high fines also face reputation damage in the labor market, estimating that high-fine plants face a 4.4% pay premium for new hires. I show that this effect is not explained by differences in worker talent using a worker effects model following Abowd et al. (1999). Overall, this evidence supports the prediction of the conceptual framework by showing that low fines are associated with hazards that are more likely to be privately efficient to remediate, while high fines are associated with cases of higher social efficiency and lower private efficiency.

I also examine the effects of regulatory enforcement on average worker earnings[1]. The theory of compensating differentials (Rosen 1974, 1986) provides testable predictions for this outcome, predicting that improved worker safety should reduce compensation. In contrast with this theory, I detect no effect of safety enforcement on earnings. This result holds for workers at both low- and high-fine plants.

The absence of wage adjustments is consistent with prior work. The evidence for adjustments to safety conditions is concentrated in findings on highly salient changes in mortality (Lavetti 2020; Lee and Taylor 2019). Safety improvements that are general or abstract are insufficiently salient to change worker beliefs. Another explanation could be that that plants rarely respond to shocks through adjustment along the wage margin (Shimer 2005). Wages may be too downward-rigid to adjust to safety changes when opposed by frictions like implicit contracts (Beaudry and DiNardo 1991) and menu costs. I contribute to this literature by identifying where and for what workers we should most expect to find compensating effects.

I also contribute to the literature on the effects of regulation through analysis of new outcomes at the plant level. Most research on regulatory effects examines firms or markets, but the effects on plants are less understood and could contain important heterogeneity. At large public firms, regulatory fines have direct effects, decreasing market values, although usually commensurately with the size of the fine (Karpoff et al. 2005). In county-level markets, the effects of regulation are more complicated, as greater stringency is associated with increased firm exit and weakened competition (Ryan 2012; Curtis 2020), suggesting effects beyond simply the financial burden of penalties.

Effects are more nuanced in the literature at the plant level, where enforcement has been shown to lower key characteristics like productivity (Greenstone et al. 2012; Li 2022). Enforcement can even spill over into plant behavior in areas unrelated to the target of the enforcement (Haviland et al. 2010; Bloom et al. 2025). My paper contributes to this literature by extending plant-level analyses, specifically considering the effects of enforcement on non-compliant firms and also how effects vary with the size of the related penalty.

This paper also contributes to the literature concerning OSHA. OSHA inspections themselves (independent of fines or enforcement) have a strong positive effect on worker safety outcomes (Levine et al. 2012; Lee and Taylor 2019). Furthermore, inspections that result in higher fines are associated with even greater improvements in safety (Haviland et al. 2010), identifying the greater effect of enforcement over inspection alone. Inspections alone have been shown to affect other economic outcomes, but the effect of fines remains only partially explored (Haviland et al. 2010). This paper emphasizes the role of fines in enforcement, exploring new outcomes, and improving our understanding of how enforcement of different OSHA standards can have differential effects.

[1]To my knowledge, the first exploration of enforcement on this outcome in a quasi-experimental setting.

II Conceptual Framework of Plant Safety Choice

Consider a population of identical plants facing the same set of workplace hazards, indexed by $h \in H$. Plants can choose to remediate a hazard h, in which case they incur a cost and benefit with certainty. The cost, C_h per worker, can be seen as reflecting the cost of additional equipment or trainings, and the benefit, δ_h per worker, as reflecting the reduced absenteeism and worker compensation costs that arise when hazards result in injury. Assume $C_h \geq 0$ and $\delta_h \geq 0$.

If plants choose not to remediate a hazard, they face two possible outcomes depending on whether or not they are visited by a regulatory inspector. With probability π, there is a surprise inspection, in which the regulator enforces remediation of all unremediated hazards and fines the plant F_h for each. With the remaining $1 - \pi$ probability, the plant is not inspected and its profitability is unaffected. Assume $\pi \in (0, 1)$, and $F_h > 0$.

II.A Safety & Enforcement Under Perfect Information

Profit-maximizing plants will remediate hazards if the expected benefits of doing so exceeds the expected costs. Under a strong assumption of perfect information (to be relaxed shortly), managers would know the value of C_h, δ_h, and F_h for all h, as well as π. They would then make all remediations where the certain outcome of remediation exceeds the expected outcome of not remediating, such that:

$$\delta_h - C_h \geq \pi(\delta_h - C_h - F_h)$$

This inequality can then be simplified to identify the role of each term. The value of remediation is increasing in the benefit δ_h which enters positively, but decreasing in the cost C_h which enters negatively. The value of leaving the hazard unremediated is decreasing in the expected fine F_h, as well as the probability of inspection.

$$\delta_h - C_h \geq -F_h \frac{\pi}{1 - \pi}$$

Hazards will be remediated by profit-maximizing plants in two cases. The first case is when the benefits to the plant exceed the costs outright, such that $\delta_h \geq C_h$, and remediation is always preferable even if fines or inspection probability were zero. The second case is when the balance among these factors is such that the net cost to the firm of remediation, $\delta_h - C_h$, is negative but still sufficiently high as to exceed the term on the right-hand side above, given F_h and π. In other words, there may be some cases where remediation is unprofitable for the plant in direct cost and benefit, but the expected fine if the hazard is left unremediated is sufficiently high to make ex-ante remediation preferable.

Consider the likely effects of enforced remediation on a plant in this context. Under perfect information, only unprofitable hazards would be left unremediated, specifically those for which the expected fines are sufficiently low relative to the net cost of remediation. This implies that remediation would always be unprofitable for the plant, effectively increasing labor costs per worker. Under most reasonable assumptions about the market structure, this should result in reduced employment, higher prices, reduced sales, and lower profits. In turn, this unprofitability would imply decreased survival for the plant affected.

II.B Safety & Enforcement Under Imperfect Information

Consider a case where plants do not perfectly perceive the cost or benefits of hazards, nor the fines and probability of inspection, but have beliefs about them[2]. This relaxation seems realistic, as OSHA enforces thousands of individual standards relating to thousands of potential hazards, and especially given managers with realistically limited scopes of control. Under this setting, plants choose ex-ante to make all remediations which they *perceive* to be preferable, where the previous inequality holds over their perceived values. I abstract from the belief-generating process, but it could be driven by factors such as inertia in safety choices.

A crucial difference emerges in contrast to the perfect information case: plants may now leave unremediated some hazards which are in reality outright profitable for the plant, $\delta_h \geq C_h$ due to mistaken beliefs that the opposite is true, $\tilde{C_h} > \tilde{\delta_h}$. This would occur if the related beliefs on fines and inspection probabilities make the perceived value of remediation sufficiently low relative to the expected fine.

When a regulator now inspects such a plant, and mandates remediation of a unaddressed hazards, it may be that the hazard is profitable or unprofitable for the plant, depending on the plant's beliefs. Accordingly, it could be that enforced remediation has the same negative effects as in the perfect information case (enforcing an unprofitable remediation) or it could be that it has opposite effects (enforcing a profitable remediation). If the regulator mandates a profitable remediation, the plant would spend C_h and achieve a benefit of $\delta_h > C_h$ per worker.

This effective decrease in the cost of labor should result in increased employment, lower prices, increased sales, and increased profits. Plants in such a case would accordingly see increased survival rates relative to the unaffected. This is of contingent on the fine being sufficiently low that it does not force the plant out of business outright, which seems justified

[2]This modification is similar to introducing frictions that generate sub-optimal behavior by firms like the managerial inertia described in DellaVigna and Gentzkow (2019).

in the OSHA context of relatively low fine values[3]. Imperfect information, whether derived from an imperfect belief-formation process or frictions like managerial inertia, demonstrates how different effects of regulatory inspections across plants could arise.

To ground the contrast between potentially profitable and unprofitable hazards in reality, consider the example of exposure to two chemicals commonly found on construction sites: lead and methylene chloride (MC). Protecting workers from lead poisoning rarely requires any special equipment, and usually it suffices that any contaminated clothing be cleaned and laundered before reuse (Occupational Safety and Health Administration 2024a). Protection from methylene chloride on the other hand requires investment in respiratory and air filtration systems (estimated to cost roughly $4,000-$9,000 per unit), as well as ongoing costs of monitoring and medical evaluations (Occupational Safety and Health Administration 2024b; Occupational Safety and Health Administration 2010). Clearly, the cost of remediation on a per-worker basis will be much lower for lead than MC.

A similar case holds for the benefits to the firm. Lead exposure is strongly connected to a variety of serious ailments, as well as affecting cognition and worker productivity or possibly even the families of workers through secondhand exposure (Occupational Safety and Health Administration 2025a). Preventing this helps the firm avoid significant costs in worker absence and medical expenses (via worker's compensation claims), while also avoiding lost productivity through cognition declines. In contrast, OSHA's own estimates for the benefit of full compliance with methylene chloride remediation nationwide is only 34 saved lives per year (Occupational Safety and Health Administration 2010).

II.C Optimal Fines and Profitability

The regulator in this framework has the objective of maximizing social welfare by ensuring workplaces are safe and achieves this by setting fines. I assume for simplicity that all hazards are socially beneficial to remediate, such that the expected social benefits of avoided injuries or death to workers always exceed the total social cost (this being comprised almost exclusively of the private cost to plant costs of remediation)[4]. The regulator achieves this maximization of safety through setting fines which incentivize plants to remediate hazards proactively, creating a threat which makes even unprofitable remediations preferable to the risk of inspection and penalization.

No fine should be necessary for those hazards which are not only socially beneficial but privately profitable as well, such that $\delta_h \geq C_h$, as the plant is always better off complying than risking a fine. However, there are many hazards for which the private unprofitability is

[3] The average fine being roughly $5,000, and even the 90th percentile being around $15,000.

[4] Of course in reality there may be some hazards for which the private costs exceed the social benefits.

dominated by the social benefits, so that the socially efficient outcome is still remediation. In such cases the regulator must set a fine to compel compliance.

Under perfect information, the fines required to compel compliance can be calculated:

$$F_h^* = \frac{\pi}{1-\pi}(C_h - \delta_h)$$

Fines would therefore be increasing in the costs of remediation and decreasing in the benefits, such that higher fines are associated with more unprofitable remediations.

The imperfect information case is more complicated, as the regulator knows that plants do not perfectly perceive the costs and benefits or remediation, nor the associated fines. The regulator cannot know with certainty what fine would ensure compliance without knowing the plants' beliefs. Furthermore, fines may even be necessary to compel compliance with privately profitable remediations, given the imperfect plant perceptions of those benefits.

While one solution would be to set extremely high fines in all cases, such that all reasonable beliefs preferred remediation, I constrain the regulator to avoid this unrealistic outcome. I limit fines to the maximum of the expected net social benefits of hazard remediation. This constraint is realistic, as fines are constrained in reality by such considerations, among other factors. One key such real-world constraint on fines is that exorbitant fines could deter even productive activity. When the definition of perfect compliance is not clear, and the cost of noncompliance is extremely high, even safe activity would be deterred out of fear.

The regulator in the imperfect information case therefore sets fines at the level which they believe will ensure compliance (given their beliefs over plant beliefs), but no higher than the expected social benefits. The assumption that all hazards are socially efficient to remediate therefore imposes positive fines in all cases.

Note that the regulator could also influence behavior by adjusting the probability of inspection, π, but I assume the regulator is constrained in their inspection capacity and always inspects the maximum that they are capable of, leaving this variable unchanged.

While extensions are possible to add further heterogeneity in the outcomes of regulatory enforcement, such as heterogeneous plants or risk aversion, this basic model provides a sufficient framework for interpreting the empirical results.

III Data

III.A Background on OSHA: Inspections & Fines

OSHA provides a rich context in which to study the effects of regulatory enforcement. OSHA is a United States federal regulatory agency with the mission to "assure America's workers have safe and healthful working conditions." It both sets safety standards and enforces them, identifying the methods by which firms must maintain safety and regularly inspecting workplaces for compliance with those methods. There are nearly 1,000 individual OSHA standards, and any given firm may be subject to several hundred at once.

OSHA enforces standards primarily through inspections of workplaces. To organize inspections, OSHA divides the country into 10 broad regions and then further into 85 Areas, each with its own office[5]. Each Area Office selects the plants in its jurisdiction for inspection, with two broad categories of selection criteria. "Programmed" inspections are targeted at specific industries or hazards, as identified by either National or Local Emphasis Programs, which consider priorities for enforcement at each level. Among the plants which meet the criteria for a given Program, OSHA usually then randomly assigns inspections to a subset of eligible plants, given its limited capacity. The second category of inspection is those which are initiated by some event. These could be formal complaints (lodged by a worker or concerned citizen), referrals (lodged by another Area Office, agency, or media organization), or major accidents (such as a worker fatality). While OSHA completes tens of thousands of inspections annually, their capacity limit typically means that they cover less than one percent of the over 8 million workplaces under their jurisdiction[6].

During an inspection, inspectors tour the facility (or broader work site in cases like construction) and review employer records to assess compliance with standards (Occupational Safety and Health Administration 2025b). This inspection is done alongside representatives of the firm, though normally without any advance notice. If a facility is found to be out of compliance with a standard, the inspector issues a citation for each separate violation. The inspection concludes with a conference between the employer representative and the inspector to discuss the results, after which the inspector returns to their Area Office to determine the relevant financial penalty or fine.

Fines themselves are calculated in accordance with a detailed process set out in the Occupational Safety and Health Act of 1970. This bases fines primarily on the gravity of the violation, which itself is determined by the combination of two sub-factors: the severity

[5] OSHA has direct jurisdiction over 28 states; the other 22 operate federally-approved state-run safety and health programs.

[6] US Department of Labor, Occupational Health and Safety Administration, "Commonly Used Statistics," https://www.osha.gov/oshstats/commonstats.html, accessed October 2025.

of the injury which could result from the violation and the probability that such an injury could occur. The severity is assessed across categories ranging from de minimus to high severity (i.e., death or permanent disability), while the probability takes into consideration the number of employees exposed, the frequency and proximity of the hazardous condition, and relevant trainings available. Fines are adjusted further to reflect the size of the employers business, as well as their good faith / history of previous violations. Once fines are levied, OSHA follows-up with plants to ensure remediation in a timely manner.

III.B OSHA Inspection Data

I use public data on OSHA inspections from the US Department of Labor Data Catalog. These provide detailed information about each inspection that OSHA conducts, including inspection type (Programmed, Initiated), plant characteristics (industry, union status, name and address), and outcomes of the inspection (violations, fines, employees affected, statutes violated).

Figure 1: Industry Shares of Inspections

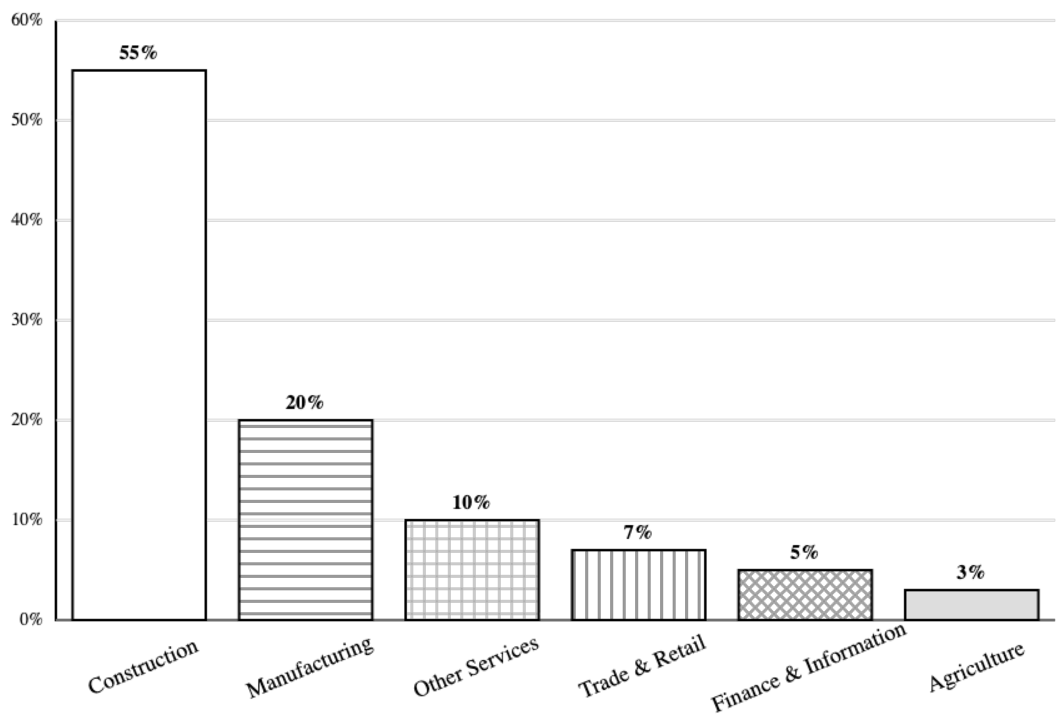

Note: The share of all plants inspected by OSHA in the sample that belong to each one-digit NAICS industry code are shown.

I restrict my sample of inspections in several ways in order to cleanly identify the effect of enforcement. To focus my analysis on inspections which are most plausibly unanticipated by the plant, I use only Programmed inspections and remove from the sample all inspections which occurred for reasons of Complaints, Referrals, or Accidents. Additionally, I use the full history of inspections to identify those plants which are experiencing their first inspection, and use only those. Finally, I restrict the sample to inspections occurring between 2000 and 2015, so that I can observe each plant for four years before and after inspection. These year cutoffs are chosen because the Census data begins in 1996 and because I want to avoid data for 2020-onwards, when the Covid-19 pandemic significantly affected OSHA operations and other economic outcomes.

Figure 1 shows a breakdown of inspections by industry. As expected, the highest shares are in those industries which include the most routine exposure to hazards, Construction and Manufacturing, accounting for three quarters of all inspections. The remaining industries are less represented, reflecting the strongly targeted nature of OSHA inspections[7]

Motivated by the conceptual framework, I analyze my sample of inspections to calculate the share of each standard across all violations, as well as the average fine-per-worker for each standard. I provide six examples below in Table 1: three associated with relatively low fines and three with relatively high fines, all of which are among the most common standards violated. These examples suggest that there is significant heterogeneity across hazards / standards.

Table 1: Example Common OSHA Violations

Violation	Share of All Violations	Average Fine / Worker
Portable Fire Extinguishers	0.5%	$50
Labeling of Hazards	0.5%	$60
Providing PPE	1%	$80
Machine Guards & Training	1%	$2,700
Eye & Face Protection	2%	$1,500
Fall Protection System	4.5%	$2,400

III.C Longitudinal Census Plant Data

To analyze long-term outcomes I link to the Longitudinal Business Database (LBD) at the US Census. The LBD is a rich annual dataset of plant-level information, covering almost all private establishments. It contains detailed firm variables, including several which are key to my analysis such as the headcount and payroll of the plant, as well as the age of the plant's broader firm. Additionally, because the LBD tracks individual establishments over time, the

[7]Mining is omitted from this analysis because of the separate agency (the Mine Safety & Health Administration) that administers safety in that industry.

closure or "death" year of a plant is identified. I use data from 1996 to 2019, allowing each plant to be observed up to four years before and after inspection. Linking the OSHA data to the LBD relies on the presence of name and street address information in both datasets. This linkage is a multi-step process, described in more detail in Appendix A. I follow best practices common in Census linking as enumerated in Chow et al. (2021).

Some outcomes of interest are not available in the LBD itself, so I additionally link from the LBD to the Longitudinal Employer-Household Dynamics (LEHD) Quarterly Workforce Indicators (QWI) dataset. This data uses administrative employer-employee histories to generate plant-level variables at the quarterly level, including hirings, separations, and earnings of both incumbent and newly hired full-time workers. I aggregate these variables to annual averages for consistency with the frequency of the LBD. The final sample includes approximately 16,400 plants[8].

IV Empirical Strategy

IV.A Propensity Score Matching of Fined & Unfined Plants

To estimate the effects of OSHA's enforcement of safety regulations, I use fines as a proxy for enforcement. The detailed fine-setting guidelines used by OSHA ensures that fines are highly reflective of the underlying danger which is being remediated through enforcement. Fines thus serve as a good measure of how intense an enforcement action is.

Comparisons of fined and unfined plants may suffer potential differences between these groups which correlate with both misbehavior and outcomes of interest. While many of these differences can be econometrically remedied through plant-level fixed effects and other controls, it is more difficult to control for differential trends between the groups. To overcome this imbalance, I use propensity score matching to create a comparison set of unfined plants which are highly similar to the fined set in both levels and trends of relevant variables. I match plants from the same state, industry, and year of inspection on the propensity to be treated given the value of key variables in the year prior to inspection. These matching variables include firm age, firm size, plant size, local HHI, plant payroll, plant average earnings, union status. For each treated fined plant, I assign a single unfined control within its state-year-industry cell, using the plant with the closest propensity score (without replacement). Within each cell, the closest matches are prioritized for producing matched pairs.

This matching yields a final analysis sample of fined plants and highly similar unfined plants evolving on similar trends pre-OSHA inspection, described in Table 2. This approach is similar to the suggestion of Abadie (2005) to use propensity score reweighting to create

[8]The Census requires rounding of key statistics such as sample sizes.

a control group with similar pre-treatment characteristics for which the parallel trends assumption is more viable. Creating an explicit control group in each state-year-industry cell and subsequently stacking the cells into one dataset also has the advantage of avoiding the complications of staggered event study timings arising from multiple times of treatment as described in Goodman-Bacon (2021). This approach of stacking matched treatment-control groups for every treatment period is similar to the related estimator proposed by Callaway and Sant'Anna (2021) for this econometric situation, as recently applied in, e.g., Schmieder et al. (2023).

Table 2: Summary Statistics

	Full Sample		Matched Sample	
	Control	Treated	Control	Treated
Firm Age	20.33	19.40	20.66	20.10
	(11.850)	(11.510)	(11.500)	(11.530)
Log Plant Size	3.11	3.15	3.15	3.16
	(1.540)	(1.453)	(1.510)	(1.420)
Log Firm Size	4.25	3.86	4.20	3.99
	(2.956)	(2.544)	(2.857)	(2.611)
Log Quarterly Worker Earnings	8.01	7.95	8.02	7.97
	(0.587)	(0.538)	(0.570)	(0.516)
Log Payroll (000s)	6.62	6.59	6.69	6.64
	(1.747)	(1.678)	(1.707)	(1.596)
Log Quarterly New Hire Earnings	7.61	7.57	7.63	7.57
	(0.696)	(0.649)	(0.686)	(0.642)
County-Sector Employment HHI	0.306	0.280	0.270	0.270
	(1.402)	(0.876)	(0.689)	(0.624)
Plants	5,900	10,500	3,600	3,600

Note: Means are provided with standard deviations in parentheses. Observations are at the plant-year, combining data from the LBD and LEHD QWI for worker earnings.

IV.B Event Study and Difference-in-Difference Specifications

My main econometric estimates are derived from the following event study specification:

$$y_{jtb} = \theta_j + \gamma_{st} + \alpha_{it} + \beta X_{jt} + \sum_{\tau=-4}^{4} \beta_\tau \mathbb{1}\{\ t = b + \tau\} + \sum_{\tau=-4}^{4} \delta_\tau \mathbb{1}\{\ t = b + \tau\}\text{Fine}_j + \epsilon_{jtb} \quad (1)$$

where y_{jtb} is the outcome of interest y for plant j in year t, and the plant was inspected by OSHA in baseline year b. I include fixed effects for the plant, θ_j, the state-year, γ_{st} and industry-year, α_{it}. X_{jt} contains time-varying plant controls, specifically firm age and HHI of local sectoral employment[9]. Finally, two sets of event study coefficients are estimated. The first set, β_τ, accounts for changes relative to inspection affecting all plants. The second set, δ_τ, are the coefficients of interest, the estimated difference between fined and unfined plants. The variable Fine_p is an indicator variable for whether plant p was fined as a result of their inspection. Each set of event study coefficients omits the period of $\tau = -1$, the year before inspection, as the reference year. Finally, ϵ_{jtb} is the error.

For every outcome of interest I also estimate the corresponding difference-in-difference model, with related specification:

$$y_{jtb} = \theta_j + \gamma_{st} + \alpha_{it} + \beta X_{jt} + \beta \mathbb{1}\{t \geq b\} + \delta \mathbb{1}\{t \geq b\}\text{Fine}_j + \epsilon_{jtb} \quad (2)$$

where the event study coefficients are simplified to β capturing the differences between pre- and post-inspection, and δ capturing the difference in the difference between fined and unfined plants pre- and post-inspection.

The key assumption for both of these models to identify the causal effect of enforcement is the conditional parallel trends assumption. This assumption means that in the absence of OSHA enforcement, the difference between the fined and unfined plants would not have changed over time. While not directly testable, the absence of differential pre-trends in the results shown in Section V support this assumption. Furthermore, timing is effectively random conditional on industry-state-year, preventing anticipation effects.

IV.C Heterogeneity by Fine Size

The conceptual framework predicts that enforcement of different OSHA standards may have different effects on plants, leading me to incorporate heterogeneity analysis. While grouping text-based standards is complex, a more natural extension is to consider the effects of variation in the size of the fine / enforcement penalty applied to plants. It is intuitive to believe

[9]The importance of labor market power in workplace conditions is shown in work like Lavetti and Schmutte (2024).

that fines of different sizes should have different effects on plants — in the extreme, it would be shocking if a fine of only a few hundred dollars was considered comparable to one in the millions.

To account for this potential heterogeneity, I conduct heterogeneity analysis by separately estimating these specifications over two sub-samples: a low-fine sample and a high-fine sample[10]. To distinguish between low and high fines, I use a threshold of $100 per worker in the fine. This roughly splits the sample into the bottom two-thirds and the top third, and is near the mean across all inspections (reflecting the upward-skew of fines generally). The per worker adjustment accounts for the fact that OSHA considers then number of employees as a key factor in fine determination. While ideally I would test multiple thresholds for robustness, the Census places strict limitations on the number of estimates which are disclosed for each project, making release of such robustness infeasible.

To consider how the effect of enforcement with low fines may differ from that of high fines, a useful lens is public efficiency in the conceptual framework of Section II. As a public regulator, OSHA seeks to minimize the distortions caused by fines in compelling the socially efficient outcome, compliance with its standards. While some remediations of hazards are privately efficient for plants to undertake, with private benefits exceeding private costs, many are likely privately inefficient (even if still socially efficient). In these cases, the private cost of remediation exceeds the private benefits to the plant but not the combined benefits to the plant and workers more broadly.

This framework predicts that OSHA should place lower fines on those standards which are more privately efficient, and higher fines on those standards which are likely privately inefficient but socially efficient. Plants are more likely to willingly remediate hazards which are privately profitable, but would require greater incentive or threat to remediate those which are privately unprofitable (even if the benefits to workers are high). The high fines on privately inefficient standards encourage plants to address hazards which benefit society despite being unprofitable for the plant.

This approach puts the most behavior-distorting fines on those standards where unregulated behavior is furthest from the social optimum, while minimizing the distortions in cases where unregulated behavior is closer to optimal. Thus, heterogeneity by fine size directly tests the model's prediction that enforcement with low fines is more likely associated with privately efficient standards (increasing plant performance), while enforcement with high fines is likely associated with privately inefficient remediations.

[10] The balance between treated and control units holds within each sub-sample, as detailed in Appendix B.

V Effects of Enforcement on Employment & Survival

V.A Low Fines Boost Employment, High Fines Lower It

I estimate Equation (1) with the outcome of log plant employment; the resulting event study coefficients are displayed in Figure 2. These results for the full sample reveal a small effect of a roughly 2% increase, which has dissipated within two years post-inspection. Applying this effect to the average plant size in this sample (50), we would expect a brief increase of about one additional worker. On its own, this result is a weak rejoinder on the broader argument that regulation depresses employment — while not a negative effect as the argument would suppose, it is only a weak (and brief) positive one.

The results by size of fine describe a more nuanced effect for regulatory enforcement. As shown in Figure 3 and Figure 4, when fines are low, plants increase employment by nearly 5% at peak, while high fines lead to similar-sized decreases in employment immediately after enforcement. In both cases, the effect appears to dissipate over a few years. Interpreting these effects for the average plant, this would imply a gain (and loss) of about two workers. The pre-treatment coefficients are all indistinguishable from zero, supporting the identification assumption.

Figure 2: Event Study ATT Estimates - Log Plant Employment

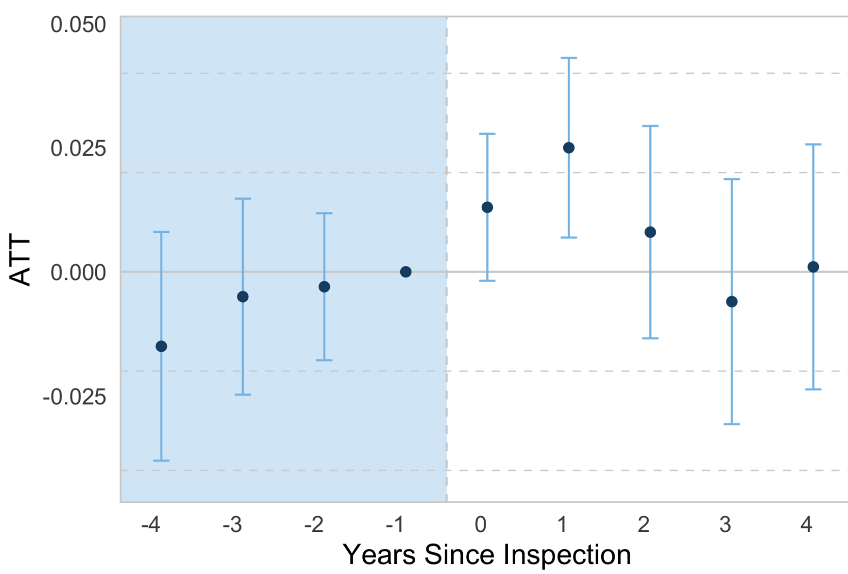

Note: Event study coefficients from Equation (1) for log plant employment relative to the inspection year (year 0), with 95% confidence intervals. The year before inspection is omitted (year -1). The shaded area represents the pre-inspection period.

Figure 3: Log Emp - Low Fine vs. None

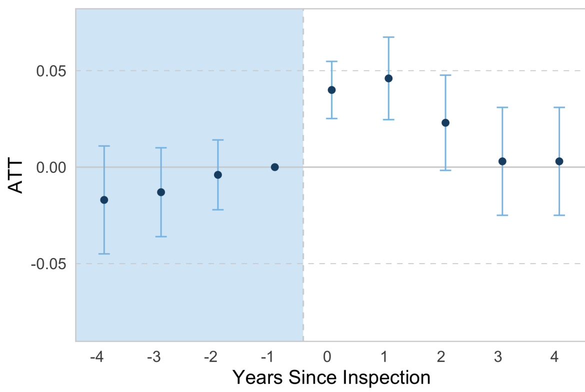

Note: Event study coefficients Equation (1) for log plant employment in the low-fine sample with 95% confidence intervals. The year before inspection is omitted (year -1).

Figure 4: Log Emp - High Fine vs. None

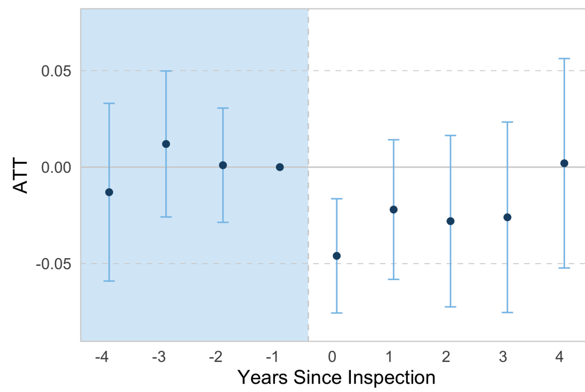

Note: Event study coefficients Equation (1) for log plant employment in the high-fine sample with 95% confidence intervals. The year before inspection is omitted (year -1).

The results of the related Equation (2) difference-in-difference models for all three samples (all, low-fine, high-fine) are reported in Table 3. While the estimated coefficients for the interaction of post-inspection and fine status are aligned in direction with the broader pattern described in the event studies, the estimates are not statistically significant, possibly due to the brief duration of the effect. This duration is unlikely to be driven by a return to noncompliance. OSHA uses prior violation record as a factor in re-inspection assignment, and the penalties for repeat violations are an order of magnitude higher than first-time.

Table 3: Difference-in-Difference Effects on Employment

	Log Plant Employment		
	(1)	(2)	(3)
	All Fines	Low Fine	High Fine
Post × Fine	0.009	0.017	-0.009
	(0.01)	(0.01)	(0.02)
	[0.366]	[0.147]	[0.630]
Plant FE	✓	✓	✓
State-Year FE	✓	✓	✓
Industry-Year FE	✓	✓	✓
R^2	0.923	0.944	0.944
N	53,500	37,000	16,500

Note: Standard errors are clustered at the matched-pair level in parentheses; P-levels in brackets.
* = p < 0.1, ** = p < 0.05, *** = p < 0.01.

V.B Low Fines Reduce Closure Probability, High Fines Raise It

The second outcome of interest is the probability of a plant closing down or going out of business. Economists usually view this outcome as a proxy for plant profitability, as plants which cannot turn a profit are those likely to close (Curtis 2020). In the LBD dataset, this is traditionally measured as the first year in which a plant reports no employment and remains without employment for additional years.

The results here are for the marginal probability of closure, i.e., the probability that a plant which was open in the previous period closes in the period of interest. Plant closure or death is more often modeled as an absorbing state, i.e., a cumulative probability rather than marginal. Presenting the results of such a model is complicated slightly by Census disclosure rules, but the results are unlikely to differ qualitatively from those presented here. In other words, these are the hazard rates rather than the cumulative failure functions.

Figure 5: Event Study ATT Estimates - Plant Marginal Closure Probability

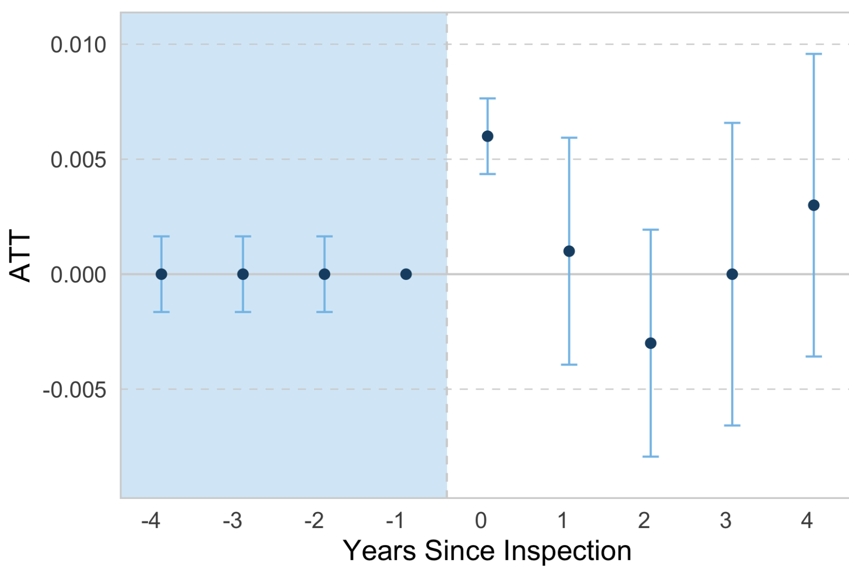

Note: Event study coefficients from Equation (1) for marginal probability of plant closure are plotted, with 95% confidence intervals. The year before inspection is omitted (year -1). The shaded area represents the pre-inspection period.

As with the outcome of employment, the results for the full sample in Figure 5 display a small, brief effect. In this case, it appears that plants which are fined are approximately 0.5 percentage points more likely to close than those which were not fined in the year of their inspection. This small absolute magnitude is large in relative terms, as only 1% of plants close in any given year on average — making the increase a 50% jump in the likelihood of closure. The estimates for years after inspection are imprecise.

Figure 6: Closure - Low Fine vs. None

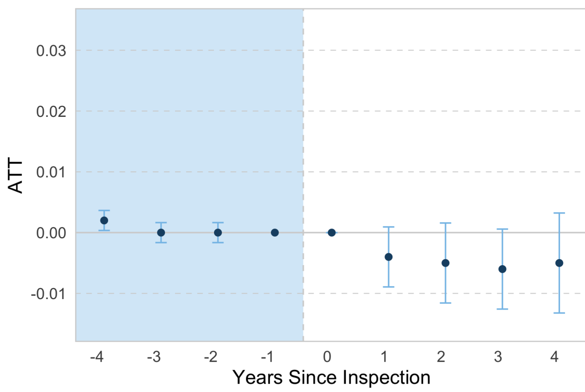

Note: Event study coefficients Equation (1) for marginal probability of plant closure in the low-fine sample with 95% confidence intervals. The year before inspection is omitted (year -1).

Figure 7: Closure - High Fine vs. None

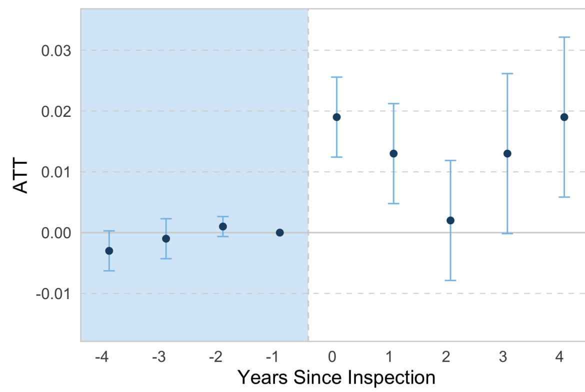

Note: Event study coefficients Equation (1) for marginal probability of plant closure in the high-fine sample with 95% confidence intervals. The year before inspection is omitted (year -1).

Table 4: Difference-in-Difference Effects on Marginal Closure Probability

	Marginal Probability of Plant Closure		
	(1)	(2)	(3)
	All Fines	Low Fine	High Fine
Post × Fine	-0.001	-0.005*	0.007
	(0.003)	(0.003)	(0.005)
	[0.591]	[0.066]	[0.159]
Plant FE	✓	✓	✓
State-Year FE	✓	✓	✓
Industry-Year FE	✓	✓	✓
R^2	0.224	0.216	0.233
N	57,500	38,500	18,500

Note: Standard errors are clustered at the matched-pair level in parentheses; P-levels in brackets.
* = p < 0.1, ** = p < 0.05, *** = p < 0.01.

Again, the simple effect displayed in the full sample belies contrasting effects when separated by low versus high fines. As shown in Figure 6, plants which receive a relatively low fine are actually less likely to close after inspection than plants that receive no fine. This result persists throughout the observation period, and is also reflected in the difference-in-difference specification, with a statistically significant effect of 0.5 percentage points in Table 4. In contrast, plants receiving high fines are not only more likely to close in the year of inspection, but also in the following years, up to the end of the observation period.

V.C Private Versus Social Efficiency Interpretation

A simple explanation of these diverging results based on the framework in Section II is that plants have imperfect information about hazards. If plants only remediate those hazards that they believe to be profitable, OSHA inspections would then enforce remediation of the remaining hazards. These could be unprofitable, increasing labor costs and lowering employment and survival, but may also be profitable, lowering labor costs and increasing employment and survival.

This interpretation of the results has a direct implication for the pattern of profitability across fines: low fines would need to be associated with profitable remediations, and high fines with unprofitable ones. This is aligned with the social efficiency lens described in Section IV.C. Low fines would be sufficient to encourage plants to remediate those hazards which are closer to privately efficient, privately profitable, while in contrast higher fines would be needed to compel compliance with those standards which are less privately efficient but more socially efficient. We would therefore predict that standards associated with low fines have greater private profitability, while standards associated with low fines have less private profitability and greater social benefits.

Assessing this relationship between fines and efficiency or profitability is difficult. While it is easy to calculate the average fine for each standard from the data, profitability is complex to measure, comprising both benefits and costs, private and social.

V.C.1 Private Versus Social Benefits

The private benefits of hazard remediation are primarily the reduced costs associated with injuries and workers' compensation, while the social benefits similarly come from avoided injuries but accrue to workers. OSHA's procedural guidelines suggest that there should be a positive relationship between both of these kinds of benefits with fines, as the primary factor in determining fines is the gravity of the potential harm resulting from the hazard, including the severity and probability (Occupational Safety and Health Administration 2025b). Thus, higher fines should be associated with those standards that prevent greater harm to workers, and therefore, provide the greatest benefit both private and social.

A crucial caveat to this relationship is that the pass-through from harm avoided for workers (social) to avoided costs for the firm (private) is not complete. There are many injuries (including death) which have statutorily-defined maximum payments in workers' compensation cases, making the private benefit effectively capped despite increasing social benefits (Oliphant and Wagner 2012). For example, the maximum penalty for a willful violation in 2025 is $165,514, including accidents which result in death (Occupational Safety and Health Administration 2025). Compare this figure with the Department of Health and

Human Service's central estimate for the value of a statistical life: \$13.1 million (Kearsley 2024). Similarly, in the state of Massachusetts a worker who is permanently blinded is entitled to a lump-sum payment of \$184,512, likely far less than most people would demand to be compensated for such an outcome[11]. There are also cases where cause and effect are more ambiguous (e.g., long-term development of disease), and workers' compensation is not as straightforward to calculate or claim. In these cases, fines are positively correlated with social benefits but no longer increasing with private benefits. At sufficiently high levels of potential worker harm, the fines increase but the private benefit stays unchanged.

An additional (non-injury related) private benefit of remediation can be conceptualized as spillover improvements from the area of remediation into other aspects of production. There is literature evidence to support the existence of such spillovers, including specifically in safety: Haviland et al. (2010) finds that enforcement of safety standards results in improvements not only in the specific standards which were enforced, but across other aspects of safety maintenance as well.

The act of bringing plants into compliance with regulations could provide a context for plants to reform other aspects of their operation. Enforcement in this view could act as a catalyst, generating effects beyond what *ex-ante* remediation would achieve. One area where these effects manifest is the management quality at the plant. Interest in this factor stems from ongoing joint work which provides evidence of a strong relationship between plant safety and management. In Bloom et al. (2025), we use similar OSHA data in combination with surveys on management practices to show that OSHA enforcement results in improved management practices at inspected plant. Given the strong literature on the connection between management and productivity (Bloom and Van Reenen 2007; Bloom et al. 2021), this effect suggests that enforcement could have further private benefits which are relevant for firm profitability.

V.C.2 Private Versus Social Costs

Private costs of remediation are a combination of direct costs (e.g., new equipment or trainings) and indirect costs (e.g., productivity loss). In contrast to benefits, costs of remediation are not officially considered in fine determination. While a comprehensive analysis of how costs vary across OSHA's nearly 1,000 standards would be prohibitively costly to undertake, I provide directional evidence through a text analysis exercise. I scraped the text of over 200 OSHA standards from the electronic federal register, pulling those which were most readily identifiable as standalone (rather than ancillary or supplementary rules). I then processed the text using a keyword algorithm which provided estimated direct costs of implementation

[11] This is based on the 2025 maximum compensation rate, \$1,922, and the specifically enumerated multiplier of that rate for lump-sum payments related to blindness, 96 (Commonwealth of Massachusetts, Executive Office of Labor and Workforce Development 2025; Commonwealth of Massachusetts 2025).

for each standard[12].

The correlation between average fines and the text-based measure of direct remediation costs is depicted in Figure 8. The relationship appears negative, with an OLS coefficient of \$0.23 lower estimated cost per dollar increase of fine.

This relationship between direct costs and fines would contradict the argument that private profitability is higher at low fines. However, it may also connect to the spillovers argument in Section V.C.1. Higher remediation costs are often driven by capital improvements, which can increase labor demand as shown in Berman and Bui (2001). These larger, physical manifestations of safety changes could provide more salient opportunities for spillover effects; however, without measures of these spillovers available, this mechanism remains suggestive.

Figure 8: Estimated Direct Remediation Costs versus Average Fines

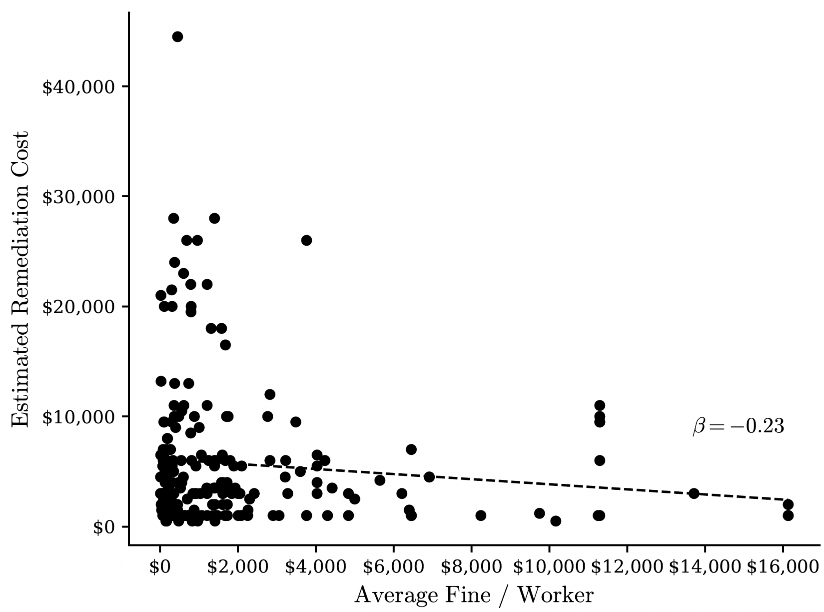

Note: Fines were calculated from the sample of all OSHA inspections. Remediation costs were estimated using a keyword algorithm.

Beyond the direct costs of compliance, a key indirect cost of remediation is productivity loss. While these are challenging to estimate based on text analysis as above, they are known to be a significant component of private compliance costs (Curtis 2020). This is especially true in the safety context. Studying the Mine Safety and Health Administration (MSHA), Li (2022) shows that enforced remediation of hazards causes productivity losses which far outweigh the benefits from reduced injury rates, particularly for the highest fines.

The MSHA uses a similar gravity-based procedure for determining fines, suggesting that

[12] The algorithm is detailed in Appendix C.

a similar effect could be observed for OSHA fines. If productivity costs are larger for high-fine violations — a plausible assumption given that severe hazards can require more disruptive operational changes as in Li (2022) — this could imply a positive relationship between fines and total private costs (direct plus indirect).

The social cost of productivity loss could also be substantial if the decrease in productivity is passed on to workers through lower wages. I consider the effect of enforcement on earnings in Section VI below, where I find no detectable decrease in average earnings for either the low- or high-fine samples. This suggests that whatever productivity losses there are from safety enforcement are impacting the plant, and not the workers.

A final cost to plants is the reputation cost of significant safety violations. Firms are known to act in response to reputation concerns relating to regulatory compliance, including regulation of wage theft (Ji and Weil 2010) and safety in particular (Johnson 2020). More broadly, Mas (2008) shows that when noncompliance is associated with worse quality, there may be effects on product markets as well. While OSHA enforcement does not necessarily result in public information about violations, OSHA publishes press releases in many cases, and they are often featured in media coverage (Johnson 2020). Furthermore, violations associated with high fines are much more likely to be publicized, both due to OSHA's own rules and public interest.

While my Census project does not allow for analysis of sales or other product market outcomes, my analysis of worker earnings provides evidence of this effect. Restricting my earnings analysis to newly hired workers, who would be choosing between plants at which to work, I find that high fines result in increased earnings costs. This suggests that high fines affect plant reputations negatively as workplaces, imposing an indirect private cost.

The weight of this evidence suggests that the relationship between fines and profitability which follows from the interpretation in Section V.C holds true. The private benefits of hazard remediation are correlated with fines, but only weakly and up to certain thresholds. Beyond this, the private benefits of remediation stall while the social benefits continue to rise with the expected damage to workers. The non-injury related benefits also accrue mainly in the case of low-fines, where literature evidence suggests that greater capital expenditure and catalyzing effects on management styles are more likely to benefit plants. On the costs side, while the direct cost of remediation may be higher for low fines, the indirect costs of productivity loss and reputation damage appear correlated with fines, suggesting that the overall private costs are much higher for higher fines. The social cost of productivity loss in contrast do not appear supported by my earnings analysis.

Overall, this evidence suggests that low fines are associated with those standards for which remediation is more privately efficient, while high fines are associated with cases of lower private efficiency, and greater social efficiency.

VI Effects of Enforcement on Worker Earnings

VI.A No Evidence of Compensating Differentials

While the argument that regulation burdens businesses does not directly imply an effect on wages, there is a strong existing literature in economics which makes predictions for the effect of safety regulation in particular. First proposed in Rosen (1974) and Rosen (1986), the theory of compensating differentials conceptualizes worker pay as reflecting not only worker contributions to production, but also their exposure to amenities. Amenities could be positive aspects of a job (e.g., free coffee) or negative ones (e.g., unsafe conditions), and affect the utility from the job accordingly. If workers are not compensated for their exposure to amenities, the resulting labor market equilibrium is inefficient: firms are paying workers more than they need to for jobs with net-positive amenities, and underpaying workers for jobs with net-negative amenities.

To apply this argument in my setting, it must be true that OSHA has positive effects on workplace safety as an amenity. OSHA has been shown to have strong positive effects on safety elsewhere in the literature. The effect is generally established in Levine et al. (2012), and further quantified in mortality terms by Lee and Taylor (2019). More closely related to my empirical design is Haviland et al. (2010), who specifically show that higher fines in OSHA inspections are associated with greater improvements in safety. Accordingly, for this paper I will assume that plants which are fined experience greater improvements in safety than those which are not fined, and that this difference is greater still for the high-fine group.

In view of this existing literature, the prediction would be that worker earnings should fall in response to OSHA enforcement events. Workers who are made safer should no longer need to be compensated as much for their previously unsafe conditions, so their earnings should fall (or not rise as quickly) relative to workers whose safety is unchanged.

I estimate a precise null effect on earnings, presented in Figure 9. The difference-in-difference version in Table 5 bounds the effect between ±1.5 percent change. Unlike in the previous two outcomes, this pattern is not overturned when analyzing the low- and high-fine groups: in both samples, no discernible effect is detected. Interpreting this result, note that my outcome of earnings masks potential offsetting effects on the underlying factors of earnings. If we simply consider earnings as hours by wage, then it could be the case that earnings remain flat while wages fall and hours increase. This would be aligned with the result of Lee and Taylor (2019), who find in a similar OSHA context a decrease in total payroll per hour worked at the plant level [13]. Indeed, an increase in hours worked in response to a safety enforcement has intuitive appeal, as safer, more cautious work is likely to be slower. Unfortunately, hours are not available for analysis under the Census project of which this paper is a part.

While a null effect is unexpected given the literature above, there are other extant arguments in the literature that could explain this result. One possibility is that the changes in safety caused by OSHA enforcement are simply too small to be salient for worker utility. Consider the foregoing example of methylene chloride — full compliance nationwide is expected to avert only 34 deaths nationwide. For the average worker, this change in expected mortality is inconceivably low, and unlikely to therefore justify a change in earnings. This aligns with previous findings from Viscusi and O'Connor (1984) that workers are not fully informed about their job hazards, and may not respond to imprecise small signals.

Figure 9: Event Study ATT Estimates - Log Average Quarterly Worker Earnings

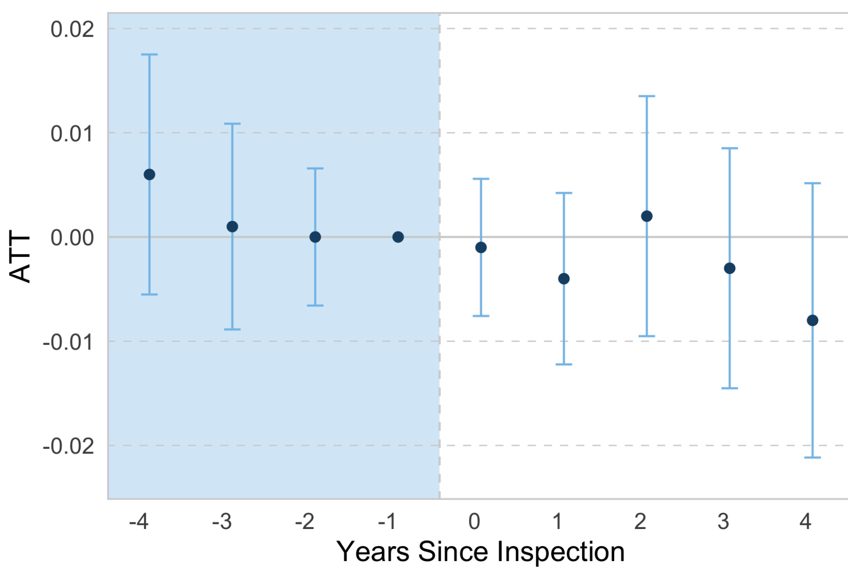

Note: Event study coefficients from Equation (1) for log quarterly full-time earnings to the inspection year (year 0), with 95% confidence intervals. The year before inspection is omitted (year -1). The shaded area represents the pre-inspection period.

The importance of amenity salience for compensating differentials to function is supported by the literature. Most of the evidence for compensating differentials is static — documenting systematic differences in earnings across occupations or industries which are explained by amenities. As stressed in Lavetti (2020)'s overview of the literature, there is scarce dynamic evidence for within-job changes in earnings which correspond to changes

[13] There are other differences between our papers which prevent a direct comparison, such as differences in industry focus and sample years.

in amenities. The quasi-experimental evidence that does exist tends to focus precisely on mortality risk as the most salient (Lee and Taylor 2019; Lavetti and Schmutte 2024).

Figure 10: ATTs - Earnings, Low Fine

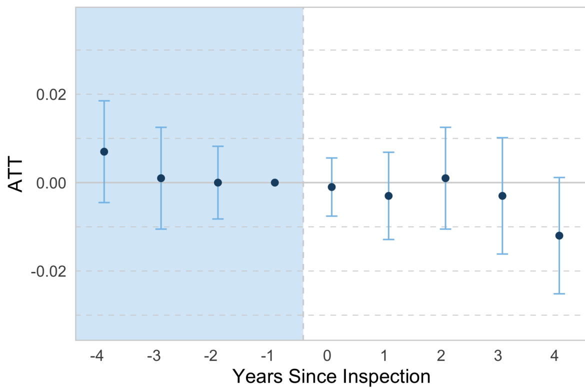

Note: Event study coefficients Equation (1) for log quarterly full-time earnings in the low-fine sample with 95% confidence intervals. The year before inspection is omitted (year -1).

Figure 11: ATTs - Earnings, High Fine

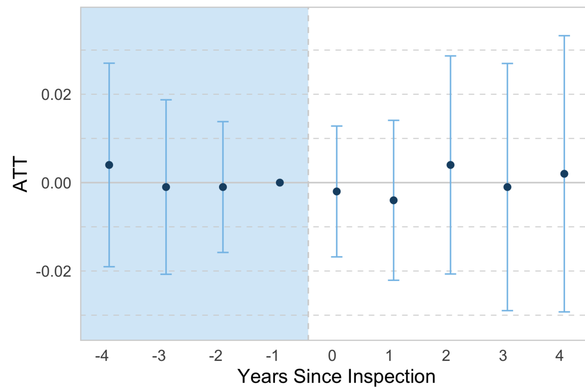

Note: Event study coefficients Equation (1) for log quarterly full-time earnings in the high-fine sample with 95% confidence intervals. The year before inspection is omitted (year -1).

Table 5: Difference-in-Difference Effects on Worker Earnings

	Log Average Quarterly Worker Earnings		
	(1)	(2)	(3)
	All Fines	Low Fine	High Fine
Post × Fine	0.000	0.000	0.000
	(0.005)	(0.003)	(0.011)
	[0.470]	[0.362]	[0.993]
Plant FE	✓	✓	✓
State-Year FE	✓	✓	✓
Industry-Year FE	✓	✓	✓
R^2	0.902	0.912	0.885
N	53,500	37,000	16,500

Note: Standard errors are clustered at the matched-pair level in parentheses; P-levels in brackets.
* = p < 0.1, ** = p < 0.05, *** = p < 0.01.

A second explanation is more general: it is possible that firms are unlikely to respond to safety changes through adjustment along the wage margin due to frictions. As originally highlighted in Shimer (2005), this concept of downward nominal wage rigidity is well established in the literature (Fallick et al. 2022; Schaefer and Singleton 2023). Even if workers did recognize the benefits of greater safety, their wages could be too downward-rigid to change due to a variety of frictions, such as implicit contracts, menu costs, and general resistance to pay cuts.

VI.B New Hires Earn More After High Fines

To test the strength of the downward rigidity explanation for null earnings results, I estimate the effect of enforcement on specifically newly hired workers. The earnings of newly hired workers are separately identified in the LEHD data, allowing me to calculate the average earnings of only newly hired full-time workers (versus all full-time workers) in each plant-year. This creates a new outcome variable, one which directly speaks to the issue of downward wage rigidities.

Even if downward wage rigidity prevents the earnings of the average full time worker from changing in response to OSHA enforcement, for newly hired workers there is no precedent wage from which to decrease. This should weaken or negate the effect of the frictions which drive the Shimer (2005) hypothesis, as firms would have greater license or freedom to set lower wages for new hires given they are only ever exposed to the new safer conditions.

Figure 12: Event Study ATT Estimates - Log Average Quarterly New-Hire Earnings

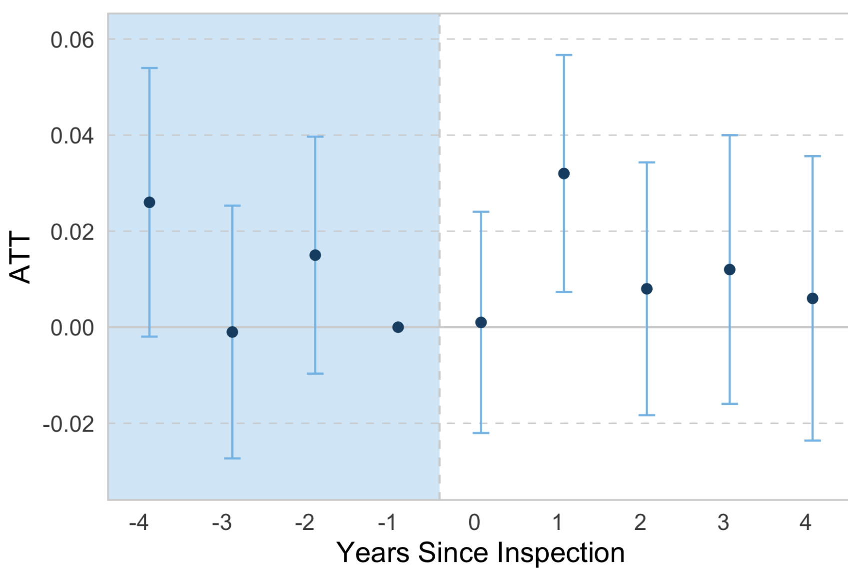

Note: Event study coefficients from Equation (1) for log quarterly new-hire earnings to the inspection year (year 0), with 95% confidence intervals. The year before inspection is omitted (year -1). The shaded area represents the pre-inspection period.

I find that OSHA enforcement does not significantly affect the earnings of newly hired workers in the full sample, as shown below in Figure 12. If anything, there is a slight increase in earnings in the first full year after the OSHA inspection takes place, of roughly 3%. Unlike earnings for all workers, the earnings of new-hires does display heterogeneity by sub-sample. Low fines display a null effect (similar to the result for all workers), while high fines surprisingly seem to have a positive effect. This is confirmed in the difference-in-difference analysis in Table 6, where an estimated effect of a 4.4% increase in earnings is estimated.

Figure 13: ATTs - New-Hire Earnings, Low Fine

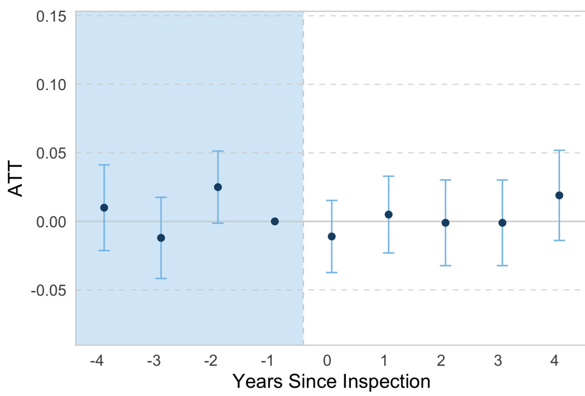

Note: Event study coefficients Equation (1) for log quarterly new-hire earnings in the low-fine sample with 95% confidence intervals. The year before inspection is omitted (year -1).

Figure 14: ATTs - New-Hire Earnings, High Fine

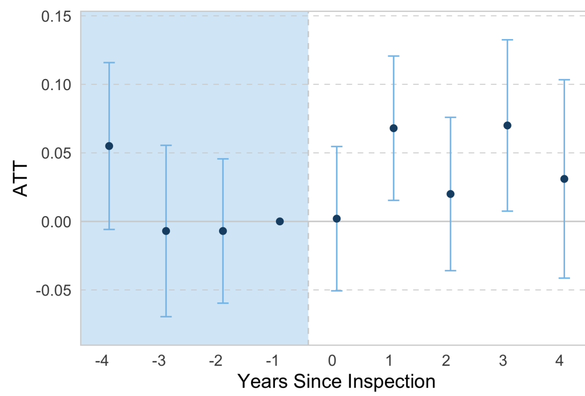

Note: Event study coefficients Equation (1) for log quarterly new-hire earnings in the high-fine sample with 95% confidence intervals. The year before inspection is omitted (year -1).

Table 6: Difference-in-Difference Effects on New-Hire Earnings

	Log Average Quarterly New-Hire Earnings		
	(1)	(2)	(3)
	All Fines	Low Fine	High Fine
Post × Fine	0.015	0.004	0.043**
	(0.009)	(0.011)	(0.020)
	[0.116]	[0.684]	[0.030]
Plant FE	✓	✓	✓
State-Year FE	✓	✓	✓
Industry-Year FE	✓	✓	✓
R^2	0.902	0.912	0.885
N	53,500	37,000	16,500

Note: Standard errors are clustered at the matched-pair level in parentheses; P-levels in brackets.
$* = p < 0.1$, $** = p < 0.05$, $*** = p < 0.01$.

The result for high fines is particularly interesting, as it is the opposite of what a compensating differentials theory would predict: safer workers are now paid more than their counterparts who did not experience an increase in safety. This contrasts not only with the prediction of compensating differentials theory, but also with the rigid wages literature. Studies of specifically newly-hired workers tend to find that they are paid quite similarly to existing workers, out of firms' desire for fairness and to avoid discouraging effort (Galuscak et al. 2012). That they are paid more in the case of high regulatory fines suggests that there is another force countervailing these frictions.

This result suggests an additional private cost to regulatory enforcement, as discussed in Section V.C.2. Enforcement of high fines may be causing reputation damage to the plant, identifying it as an unsafe workplace. It may also be a signal of a poorly managed plant, following the findings of Bloom et al. (2025). In this way, OSHA enforcement of particularly high fines may act as an information shock to potential new hires, causing them to demand higher wages as a result. This aligns with the effects found in Viscusi and O'Connor (1984) using experimental information treatments about safety. Workers could find out about OSHA enforcement through a variety of mechanisms, including press releases, media coverage, or word of mouth in particular industries.

This finding identifies a concrete private cost for firms experiencing high fines: higher labor costs for new hires. This additional cost supports the profitability-based interpretation of the results; however, this support is contingent upon the assumption that the higher wages are not reflective of efficient payment for higher talent / ability. If high-fined plants are paying more for new hires, it could be the case that they are hiring more talented workers and getting more productivity out of them, which would negate the effect on profitability.

To test this assumption, I estimate a model which identifies the average worker talent through fixed effects analysis (Abowd et al. 1999). The results are reported in full in Appendix D, but do not support the assertion that high-fine plants are gaining more talented workers after inspection. If anything, the results suggest a negative overall effect on worker talent across fined plants.

VII Conclusion

Opponents of business regulation argue it stifles economic activity, depressing employment and hurting businesses. While this may hold in some contexts (Karpoff et al. 2005; Ryan 2012; Curtis 2020), I provide new evidence that the micro-level effects of regulatory enforcement are more complicated. I find that the effects of enforcement depend heavily on the balance between private and social efficiency of remediation.

While enforcement resulting in relatively high fines can depress employment and make plants more likely to close, in contrast when relatively low fines are assessed the opposite is true: plants increase employment and are more likely to remain open. This likely reflects OSHA's approach to fines as a social regulator, seeking to minimize distortions of behavior in achieving a socially efficient outcome. In cases where standards are more privately efficient, only lower fines are required to compel compliance. In contrast, cases of lower private efficiency but greater social efficiency require higher fines, in order to compel firms to comply despite the high private costs such as productivity loss and reputation damage.

These findings could inform improvements in regulatory design. OSHA inspectors could more proactively communicate about the opportunity for spillover improvements, building on their experience across plants which have previously remediated similar hazards successfully. They could also help mitigate indirect costs, by emphasizing safety improvements to the public when they occur, potentially reducing the reputation damage of high fines.

This analysis holds for safety regulation, but may not be applicable to other areas of regulation and regulators. Safety remediation has relatively clear implications in cost-saving for the firm, whereas environmental or workers-rights regulation does not. Additionally, the effects of OSHA enforcement may be driven by agency-specific nuances; future research into cross-regulator differences would help understand the external validity of these results.

Regulatory enforcement does not have uniform effects. They depend on what standards are being enforced, the associated penalties, and the costs and benefits which come from remediation. This finding encourages future research into the multidimensional effects of regulatory enforcement, to better understand the indirect impacts on plants. Additionally, studies which can estimate the long-term impacts on workers would provide a more holistic understanding of the welfare effects of regulatory enforcement.

References

Abadie, A. (2005). Semiparametric difference-in-differences estimators. *The Review of Economic Studies*, 72(1):1–19.

Abowd, J. M., Kramarz, F., and Margolis, D. N. (1999). High wage workers and high wage firms. *Econometrica*, 67(2):251–333.

Alfaro-Urena, A., Faber, B., Gaubert, C., Manelici, I., and Vasquez, J. (2022). Responsible Sourcing? Theory and Evidence from Costa Rica. Working Paper 30683, National Bureau of Economic Research.

Beaudry, P. and DiNardo, J. (1991). The effect of implicit contracts on the movement of wages over the business cycle: Evidence from micro data. *Journal of Political Economy*, 99(4):665–688.

Berman, E. and Bui, L. T. (2001). Environmental regulation and labor demand: evidence from the South Coast Air Basin. *Journal of Public Economics*, 79(2):265–295.

Bloom, N., Kawakubo, T., Meng, C., Mizen, P., Riley, R., Senga, T., and Van Reenen, J. (2021). Do well managed firms make better forecasts? Working Paper 29591, National Bureau of Economic Research.

Bloom, N., Levine, D., Johnson, M., and Watson, T. (2025). Safety and management. *Working Paper*.

Bloom, N. and Van Reenen, J. (2007). Measuring and explaining management practices across firms and countries. *The Quarterly Journal of Economics*, 122(4):1351–1408.

Branstetter, L., Lima, F., Taylor, L. J., and Venâncio, A. (2013). Do entry regulations deter entrepreneurship and job creation? evidence from recent reforms in portugal. *The Economic Journal*, 124(577):805–832.

Callaway, B. and Sant'Anna, P. H. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2):200–230. Themed Issue: Treatment Effect 1.

Card, D., Heining, J., and Kline, P. (2013). Workplace heterogeneity and the rise of west german wage inequality. *The Quarterly Journal of Economics*, 128(3):967–1015.

Chow, M., Fort, T., Goetz, C., Goldschlag, N., Lawrence, J., Perlman, E., Stinson, M., and White, T. K. (2021). Redesigning the longitudinal business database. Working Paper 21-08, Center for Economic Studies.

Commonwealth of Massachusetts (2025). Massachusetts General Laws, Part I, Title

XXI, Chapter 152, Section 36: Specific Injuries. https://malegislature.gov/Laws/GeneralLaws/PartI/TitleXXI/Chapter152/Section36. Accessed November 3, 2025.

Commonwealth of Massachusetts, Executive Office of Labor and Workforce Development (2025). Minimum and Maximum Compensation Rates. https://www.mass.gov/info-details/minimum-and-maximum-compensation-rates. Accessed November 3, 2025.

Curtis, E. M. (2020). Reevaluating the ozone nonattainment standards: Evidence from the 2004 expansion. *Journal of Environmental Economics and Management*, 99:102261.

DellaVigna, S. and Gentzkow, M. (2019). Uniform pricing in u.s. retail chains. *The Quarterly Journal of Economics*, 134(4):2011–2084.

Fallick, B., Villar, D., and Wascher, W. (2022). Downward nominal wage rigidity in the united states in times of economic distress and low inflation. *Labour Economics*, 78:102246.

Galuscak, K., Keeney, M., Nicolitsas, D., Smets, F., Strzelecki, P., and Vodopivec, M. (2012). The determination of wages of newly hired employees: Survey evidence on internal versus external factors. *Labour Economics*, 19(5):802–812. Special Section on: Price, Wage and Employment Adjustments in 2007-2008 and Some Inferences for the Current European Crisis.

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, 225(2):254–277. Themed Issue: Treatment Effect 1.

Greenstone, M., List, J. A., and Syverson, C. (2012). The effects of environmental regulation on the competitiveness of u.s. manufacturing. Working Paper 18392, National Bureau of Economic Research.

Haviland, A., Burns, R., Gray, W., Ruder, T., and Mendeloff, J. (2010). What kinds of injuries do OSHA inspections prevent? *Journal of Safety Research*, 41(4):339–345.

Ji, M. and Weil, D. (2010). Does ownership structure influence regulatory behavior? the impact of franchising on labor standards compliance. *Industrial and Labor Relations Review*, 68(5):977–1006.

Johnson, M. S. (2020). Regulation by shaming: Deterrence effects of publicizing violations of workplace safety and health laws. *American Economic Review*, 110(6):1866–1904.

Karpoff, J., John R. Lott, , and Wehrly, E. (2005). The reputational penalties for environmental violations: Empirical evidence. *The Journal of Law Economics*, 48(2):653–675.

Kearsley, A. (2024). HHS Standard Values for Regulatory Analysis. https://aspe.hhs.gov/sites/default/files/documents/cd2a1348ea0777b1aa918089e4965b8c/

standard-ria-values.pdf. ASPE Data Point, U.S. Department of Health & Human Services. Accessed November 3, 2025.

Lavetti, K. (2020). The Estimation of Compensating Wage Differentials: Lessons From the Deadliest Catch. *Journal of Business & Economic Statistics*, 38(1):165–182.

Lavetti, K. and Schmutte, I. (2024). Estimating compensating wage differentials with endogenous job mobility. *Working Paper*.

Lee, J. M. and Taylor, L. O. (2019). Randomized safety inspections and risk exposure on the job: Quasi-experimental estimates of the value of a statistical life. *American Economic Journal: Economic Policy*, 11(4):350–74.

Levine, D. I., Toffel, M. W., and Johnson, M. S. (2012). Randomized government safety inspections reduce worker injuries with no detectable job loss. *Science*, 336(6083):907–911.

Li, L. (2022). Workplace Safety and Worker Productivity: Evidence from the MINER Act. *ILR Review*, 75(1):117–138.

Mas, A. (2008). Labour unrest and the quality of production: Evidence from the construction equipment resale market. *The Review of Economic Studies*, 75(1):229–258.

Occupational Safety and Health Administration (2024a). Lead. 29 C.F.R. § 1910.1025. US Department of Labor.

Occupational Safety and Health Administration (2024b). Methylene chloride. 29 C.F.R. § 1910.1052. US Department of Labor.

Occupational Safety and Health Administration (2025). 2025 annual adjustments to osha civil penalties. Memorandum, U.S. Department of Labor. Accessed November 3, 2025.

Occupational Safety and Health Administration (2010). Regulatory review of 29 CFR 1910.1052: Methylene chloride pursuant to Section 610 of the Regulatory Flexibility Act and Section 5 of Executive Order 12866. Regulatory review, US Department of Labor.

Occupational Safety and Health Administration (2025a). Lead: Health effects. Technical report, US Department of Labor.

Occupational Safety and Health Administration (2025b). OSHA's Field Operation Manual (FOM), Directive Number CPL 02-00-148.

Oliphant, K. and Wagner, G. (2012). *Employers' Liability and Workers' Compensation*. Tort and Insurance Law. De Gruyter.

Rosen, S. (1974). Hedonic prices and implicit markets: Product differentiation in pure competition. *Journal of Political Economy*, 82(1):34–55.

Rosen, S. (1986). Chapter 12: The theory of equalizing differences. Handbook of Labor Economics, pages 641–692. Elsevier.

Ryan, S. P. (2012). The costs of environmental regulation in a concentrated industry. *Econometrica*, 80(3):1019–1061.

Schaefer, D. and Singleton, C. (2023). The extent of downward nominal wage rigidity: New evidence from payroll data. *Review of Economic Dynamics*, 51:60–76.

Schmieder, J. F., von Wachter, T., and Heining, J. (2023). The Costs of Job Displacement over the Business Cycle and Its Sources: Evidence from Germany. *American Economic Review*, 113(5):1208–54.

Shimer, R. (2005). The cyclical behavior of equilibrium unemployment and vacancies. *American Economic Review*, 95(1):25–49.

Stansbury, A. (2021). Do US firms have an incentive to comply with the FLSA and the NLRA? Working Paper 21-9, Peterson Institute for International Economics.

US Chamber of Commerce (2025). Overregulation Is Crippling Business, Getting Regulations Right Is the Key to Growth.

Viscusi, W. K. and O'Connor, C. J. (1984). Adaptive responses to chemical labeling: Are workers bayesian decision makers? *The American Economic Review*, 74(5):942–956.

Appendix

A Data Linkages

This section describes the two major data linkages which were required to conduct the empirical analysis of this project. The first linkage was between the public OSHA inspection archive and the LBD at the US Census. The second was from the LBD to the LEHD within the Census.

The LBD is often linked to data external to the Census via an intermediary dataset, the County Business Patterns Business Register (CBPBR). This dataset operates similarly to the LBD, including information from every business in every year, but crucially has the additional variables of name and address (whereas the LBD only includes anonymized identifier variables, such as LBD Number). Individual plant-year observations are cleanly linked one-to-one between the LBD and CBPBR using common administrative variables in both, called `id` in the CBPBR and `estabid` in the LBD.

Once the LBD is linked to the CBPBR, the names and addresses of each plant can be used in a fuzzy-linking methodology to link outwards to the OSHA data. Helpfully, both datasets contain not one but two name variables, as is common practice for businesses. Usually, these reflect one "trade name" and one more formal "doing-business-as" name (for example, "Restaurant Holdings LLC" and "John's Restaurant"). The following procedure was used:

1. Block data into cells of observation year, state, county, and industry in both datasets
2. Within each cell, conduct fuzzy matching of observations between the datasets on address and name variables, using the `reclink2` command in Stata (conducting a separate matching for each of the four possible name variable permutations)
3. Refine the identified viable matches to restrict to the one OSHA observation with the highest matching score (produced by `reclink2` for string similarity) for each LBD observation
4. Use simple string processing rules to eliminate matches which are clearly poor quality
5. Spot check the remaining matches for one-off corrections

The resulting dataset matched approximately 40% of the OSHA archive into the LBD, on par with similar efforts in related projects (Bloom et al. 2025).

The second linking was internal to the Census, between the LBD and the LEHD QWI dataset. Unlike the LBD and CBPBR, there is no single common identifying variable between these two datasets at the plant-year observation level. Instead, there is a variable which is common to both and identifies unique firms in each year, called `EIN` (for Employer Identification Number) in the LBD and `SEIN` in the LEHD (State EIN). While it is possible to match observations at the plant-year level one-to-one using this variable, that is only true for those plants which are the sole operating plant of the broader firm in a given state. In other words, Firm A can have one plant in Maine and two plants in Massachusetts. Using this method, it would be possible to link the plant in Maine as only one plant would appear in Maine with Firm A's EIN in both datasets. For the two plants in Massachusetts, both appear with the same EIN in both datasets, making the matching of which-is-which more difficult.

To overcome this difficulty, I leverage additional variables which are similar in both datasets to try and refine the matching. Specifically, I use the industry codes and detailed geographic variables to restrict the number of possible matches to one in each case. For example, if the two plants in Massachusetts above are in different counties, I would then match between the two datasets on year-EIN-county, and find only one possible match in each case. Alternatively, if the two plants are in the same county but one is a manufacturing facility while the other is a logistics trade hub, they would have different industry codes, but the same codes in each dataset. I iterate over several combinations of such identifying variables to identify all possible unambiguous one-to-one matches between the two datasets.

B Balance Table

The balance between treatment and control samples is also balanced within the two sub-samples defined by the low versus high fine treatments.

Table 7: Summary Statistics

	Low Fine		High Fine	
	Control	Treated	Control	Treated
Firm Age	20.9	21.0	20.1	18.2
	(11.090)	(11.000)	(12.290)	(12.360)
Log Plant Size	3.32	3.60	2.79	2.16
	(1.512)	(1.305)	(1.441)	(1.131)
Log Firm Size	4.50	4.61	3.58	2.71
	(2.877)	(2.601)	(2.713)	(2.119)
Log Quarterly Worker Earnings	8.01	7.98	8.05	7.94
	(0.553)	(0.480)	(0.604)	(0.584)
Log Payroll (000s)	6.84	7.10	6.36	5.61
	(1.692)	(1.455)	(1.692)	(1.405)
Log Quarterly New Hire Earnings	7.61	7.58	7.68	7.55
	(0.682)	(0.611)	(0.695)	(0.716)
County-Sector Employment HHI	0.283	0.299	0.244	0.209
	(0.699)	(0.678)	(0.668)	(0.484)
Plants	2,600	2,600	1,200	1,200

Note: Means are provided with standard deviations in parentheses.

C Estimating OSHA Standard Remediation Costs

I used a keyword-based text algorithm to estimate the costs of compliance with OSHA standards. This approach combines tractability and replicability with sufficient richness to generate variation.

I constructed a dataset of OSHA texts by scraping the Electronic Code of the Federal Register, using the standards enumerated in the OSHA inspections data. I scraped the specific text of the regulatory statute — for example, for standard 1926.1153(g)(4), I extract:

"The employer shall designate a competent person to make frequent and regular inspections of job sites, materials, and equipment to implement the written exposure control plan."

I then studied the key words which appeared frequently across standards and developed a set of cost logic rules for the algorithm, displayed in Table 8. Costs were based on an assumption of 10 employees.

Table 8: Cost Estimation Rules Based on Standard Keywords

Category	**Keywords**	**Cost Logic**
Training Costs	train, training program, ensure employee, inform employees	+\$1,500 +\$500 if annual
Medical Surveillance	medical, healthcare professional, physician, examination, medical surveillance	+\$3000; +\$1000 if annual or periodic
Respiratory Protection Program	respirat	+\$6,000
Exposure Monitoring	monitor, monitoring, exposure assessment, air sample, breathing zone	+\$4,000; +\$2,000 if initial and periodic; +\$3,000 if quarterly
Written Program Development	written program, written plan, exposure control plan, establish a written	+\$2,000; +\$500 if review and update
Personal Protective Equipment	protective clothing, protective equipment, ppe, gloves, masks, face shield, goggles, apron, gown, coverall	+\$1,000; +\$500 if provide and no cost
Engineering Controls	engineering control, ventilation, enclosure, exhaust, isolation, local exhaust	+\$20,000; +\$5,000 if hepa
Signage and Labeling	sign, label, posted, warning, caution, danger	+\$500
Record Keeping	record, maintain records, documentation, keep a copy	+\$1,200
Inspection Programs	inspect, inspection, qualified person	+\$2,000; +\$1,500 if monthly; +\$2,500 if daily or each shift
Specialized Equipment	ladder, scaffold, crane, hoist, rigging	+\$3,000
Hazard Communication Program	hazard communication, safety data sheet, sds	+\$3,000
Emergency Procedures	emergency, evacuation, rescue, emergency response	+\$3,500
Confined Space Program	permit space, confined space	+\$6,000
Blood	blood, pathogen, hepatitis, hiv, sharps	+\$5,000
Asbestos	asbestos, acm	+\$5,000
Toxic Chemicals	lead, silica, formaldehyde, cadmium, chromium	+\$5,500
Facilities	shower, change room, lunchroom, washing facilities	+\$10,000
Competent Person Designation	competent person	+\$5,000
Cleaning and Housekeeping	clean, housekeeping, decontamination, waste disposal	+\$2,000

D AKM Analysis

When estimating the effects of enforcement on worker earnings, analysis at the plant level suffers from a lack of information about the underlying workers. Without knowing the distribution of worker characteristics (particularly worker ability), the interpretation of plant-level regressions is limited. Changes could be efficient responses to changes in workers, or could be inefficient changes resulting from frictions.

For example, if wages increase following enforcement, this could reflect of increasing worker talent or ability leading to efficiently higher wages Alternatively, it could be that the plant is forced to overpay relative to talent due to reputational damage from OSHA enforcement (with no change in worker ability). Without worker information, we cannot disentangle these two explanations.

I estimate a model of worker earnings following Abowd et al. (1999), a common way to overcome this problem. This model (henceforth AKM) incorporates both worker fixed effects and plant fixed effects to the earnings process. This allow me to then consider the estimated worker fixed effects as a proxy for worker ability, as it measures how much individuals earn relative to expectations based on their places of employment. I use an AKM model to estimate the average worker earning ability at each plant in each year, using the following two-step approach. In a first step, I estimate a simple AKM model using all workers, plants, and years available in the LEHD data:

$$\ln w_{it} = \omega_i + \pi_{j(it)} + \beta X_{it} + \epsilon_{it} \tag{3}$$

where w_{it} is the total earnings of worker i in quarter t[14]. This is modelled as a function of the worker's own fixed effect, ω_i, the effect of the plant at which the worker is in that quarter, $\pi_{j(it)}$, and a vector of time varying individual control variables X_{it}, which is essentially age polynomials interacted with education following Card et al. (2013). Although it is possible in principle to have estimated this step using only the analysis sample, the smaller sample size would have meant significantly noisier estimates of the effects.

In a second step, I then estimate the average $\hat{\omega}_i$ of workers in set I at each plant j in each year t as:

$$\bar{\omega}_{jt} = \frac{\sum_{i \in I_{jt}} \hat{\omega}_i}{I_{jt}} \tag{4}$$

I then use this $\bar{\omega}_{jt}$ as an outcome in the aforementioned event study and difference-in-

[14]The LEHD identifies all jobs at which a worker is employed in every quarter, resulting in multiple earnings streams. I restrict the sample to include only one observation per worker-quarter, using that of the highest earnings value.

difference specifications. The results presented in Figure 15 suggest that worker talent is not significantly impacted by OSHA enforcement. The point estimates suggest a possible immediate rise in average worker talent post enforcement, followed by a steady decline in the subsequent years, though the effects are small and all estimates are imprecise.

Figure 15: Event Study ATT Estimates - Average Worker Effect

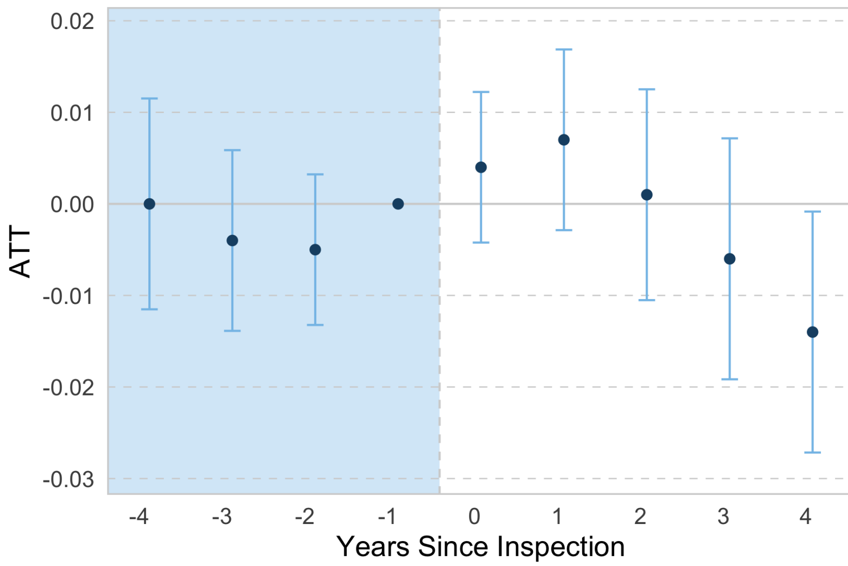

Note: Event study coefficients from Equation (1) for average worker effects relative to the inspection year (year 0), with 95% confidence intervals. The year before inspection is omitted (year -1). The shaded area represents the pre-inspection period.