

The Balloon That Didn't Inflate: Testing Cross-Hazard Injury Substitution from OSHA's Silica Standard

APEP Autonomous Research* @ailscl

March 27, 2026

Abstract

When regulators target a specific workplace hazard, do firms squeeze their finite safety budgets and inflate risk elsewhere? I test this “safety balloon” hypothesis using OSHA’s crystalline silica standard and 1.3 million establishment-hazard-year observations from the Injury Tracking Application. A triple-difference design compares targeted (respiratory) and non-targeted (injuries, hearing loss, skin disorders) hazard categories within high-silica versus low-silica manufacturers, before and after the standard’s engineering controls deadline. The balloon did not inflate. Among illness categories alone, the substitution estimate is a precise zero (+0.002, SE = 0.003). Including total injuries, the triple-difference turns *negative* (−0.032, $p = 0.002$), driven by a large decline in physical injuries at high-silica establishments—consistent with engineering controls (enclosed workstations, ventilation) reducing both dust exposure and traumatic injury risk. However, this result attenuates substantially when excluding COVID-era years, cautioning against strong causal claims.

JEL Codes: J28, K32, L51

Keywords: workplace safety, OSHA, silica standard, regulatory substitution, cross-hazard spillovers

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 1h 1m).

1. Introduction

A factory installs new ventilation systems to meet a dust-exposure limit. The engineering overhaul consumes the year’s safety budget. Meanwhile, a guard rail goes unrepaired and a hearing-protection program is deferred. Does the targeted regulation make the workplace safer on balance, or merely rearrange which workers get hurt?

This concern—that hazard-specific regulation creates a “safety balloon,” squeezing risk in one dimension only to inflate it in another—is a first-order question for regulatory design. It connects to the broader Peltzman effect in safety regulation (Peltzman, 1975; Viscusi, 1984): if agents respond to risk reduction on one margin by taking more risk on another, the net welfare gain from targeted mandates may be substantially overstated. Yet the hypothesis has never been tested in the workplace safety literature. The extensive body of work evaluating OSHA enforcement examines total injury rates (Viscusi, 1979; Bartel and Thomas, 1985; Gray and Scholz, 2005; Levine et al., 2012; Johnson, 2020; Johnson et al., 2023), without asking whether improvements in targeted hazard categories come at the expense of others.

This paper provides the first direct test. I exploit OSHA’s 2018 respirable crystalline silica standard for general industry—which reduced the permissible exposure limit (PEL) from 250 to 50 $\mu\text{g}/\text{m}^3$, imposing substantial compliance costs on silica-intensive manufacturers (OSHA, 2016)—and pair it with OSHA’s Injury Tracking Application (ITA), which provides establishment-level injury and illness counts disaggregated by hazard category for over 43,000 manufacturing establishments from 2016 to 2024.

The identification strategy is a triple-difference. I compare: (1) high-silica subsectors (NAICS 327 stone/glass/cement, 331 primary metals, 332 fabricated metals) versus low-silica manufacturing; (2) before versus after the standard’s June 2021 engineering controls deadline—the most expensive compliance phase; and (3) targeted hazard categories (respiratory conditions) versus non-targeted categories (injuries, hearing loss, skin disorders, other illnesses) within the same establishment and year. The saturated specification includes establishment-by-hazard, hazard-by-year, and establishment-by-year fixed effects, isolating the differential change in non-targeted hazards relative to targeted hazards at high-silica establishments—the substitution margin.

The results reject the safety-balloon hypothesis decisively—in the wrong direction for substitution. The saturated triple-difference coefficient is $\hat{\beta}_3 = -0.032$ (SE = 0.010, $p = 0.002$), meaning non-targeted hazards *improved more* than targeted hazards at high-silica establishments after the engineering controls deadline. Total injury rates also fell significantly: -0.135 per 100 FTE ($p = 0.002$, SDE = -0.03). Rather than substituting across hazard categories, firms that invested in silica compliance improved their overall safety performance.

Two features make this finding informative. First, the power is substantial: 1.3 million establishment-hazard-year observations, 11,545 treated establishments, and six pre-treatment years (2016–2021) that span the entire regulatory ramp-up. Second, the result is robust: broadening the silica definition to include NAICS 321 and 324 strengthens the coefficient (-0.052 , $p < 0.001$); narrowing to NAICS 327 alone yields -0.049 ($p = 0.015$).

The event study reveals the phased implementation timeline. Prior to the 2018 PEL effective date ($t - 5$ to $t - 4$), the triple-difference is near zero. During the PEL phase-in ($t - 3$ to $t - 2$, covering 2018–2019), the coefficient becomes increasingly negative, consistent with firms beginning to adjust safety investments. The pre-trend in these years reflects the gradual compliance ramp-up, not a violation of the identifying assumption. After full implementation ($t + 1$ onward), the coefficient reaches -0.126 by 2024, indicating accelerating complementarity between hazard-specific compliance and general workplace safety.

This paper contributes to three literatures. First, within occupational safety, it addresses the gap identified by [Shimshack \(2014\)](#) and [Johnson et al. \(2023\)](#): understanding how firms allocate safety investments across margins when facing targeted regulatory pressure. My null finding suggests that compliance with the silica standard either drew on slack resources or generated positive spillovers to general safety culture, rather than crowding out other protections. Second, the paper extends the Peltzman-effect tradition ([Peltzman, 1975](#); [Viscusi, 1992](#)) to a new domain with an unusually clean test: within-establishment, within-year variation across hazard categories, where the “moral hazard” channel should be minimal. Third, methodologically, the paper demonstrates the value of the ITA’s hazard-category disaggregation for regulatory evaluation—a feature of these data that has been largely unexploited.

The rest of the paper proceeds as follows. [Section 2](#) describes the silica standard and the ITA data. [Section 3](#) presents the sample and summary statistics. [Section 4](#) details the triple-difference design. [Section 5](#) reports the main results, event study, and robustness. [Section 6](#) discusses mechanisms and policy implications. [Section 7](#) concludes.

2. Institutional Background

The Crystalline Silica Standard. Respirable crystalline silica is generated when workers cut, grind, or drill materials containing quartz—principally in construction, stone fabrication, foundries, glass manufacturing, and cement production. Chronic exposure causes silicosis, lung cancer, and chronic obstructive pulmonary disease. OSHA’s previous PEL for general industry, set in 1971, was $250 \mu\text{g}/\text{m}^3$ —five times the level that the scientific literature had established as harmful by the 1990s.

OSHA published its final rule on March 25, 2016, reducing the PEL to 50 $\mu\text{g}/\text{m}^3$ for all industries (OSHA, 2016). The compliance timeline was staggered: construction firms faced the new PEL beginning June 23, 2016, while general industry (including manufacturing) had until June 23, 2018 for the PEL and until June 23, 2021 for engineering controls and medical surveillance. The standard requires employers to: measure workplace silica exposure, implement engineering controls (ventilation, wet cutting), provide respiratory protection, conduct medical surveillance for exposed workers, and maintain exposure records. OSHA estimated annualized compliance costs of \$637 million for general industry, with costs concentrated in stone/clay/glass (NAICS 327), primary metals (NAICS 331), and fabricated metals (NAICS 332).

The Injury Tracking Application. Since 2016, OSHA’s electronic reporting rule has required covered establishments to submit Form 300A data annually. These data record: total injuries, total skin disorders, respiratory conditions, poisonings, hearing loss, and other illnesses, along with days away from work (DAFW), days of job transfer/restriction, total deaths, average employment, and total hours worked. The establishment identifier enables longitudinal tracking. Reporting requirements changed over the sample period: in 2016–2018, both large establishments (250+ employees, all industries) and smaller establishments (20–249 employees, designated high-hazard industries) reported. In 2019, the smaller-establishment requirement was temporarily removed. In 2023, establishments with 100+ employees in high-hazard industries were again required to report.

3. Data

I download all ITA Form 300A summary files from OSHA’s website for reporting years 2016–2024, yielding 2.8 million establishment-year observations. I restrict to manufacturing (NAICS 31–33), drop establishments with missing or zero employees or hours, and construct a panel of establishments present for at least four years, producing 254,933 establishment-years from 43,954 unique establishments.

I classify establishments by silica exposure intensity based on their 3-digit NAICS code. *High-silica* subsectors—NAICS 327 (nonmetallic mineral products), 331 (primary metals), and 332 (fabricated metal products)—have documented high silica exposure from grinding, cutting, and foundry operations. All remaining manufacturing subsectors serve as the *low-silica* comparison group. The treatment date is 2022, the first full calendar year after the June 2021 engineering controls deadline—the most expensive compliance phase.

Injury and illness rates are expressed per 100 full-time equivalent (FTE) workers, computed

Table 1: Summary Statistics: Manufacturing Establishments, Pre-Treatment (2016–2021)

	High-Silica NAICS 327/331/332	Low-Silica Other Manufacturing
Establishments	8,597	23,683
Establishment-years	19,080	53,822
Mean employees	102	227
<i>Rates per 100 FTE-years</i>		
Total injuries	27.38	6.32
Respiratory conditions	0.017	0.018
Hearing loss	0.215	0.109
Skin disorders	0.057	0.042
Other illnesses	0.098	0.322

Notes: Rates computed as (count / total hours worked) \times 200,000. High-silica subsectors: NAICS 327 (stone, clay, glass, cement), 331 (primary metals), 332 (fabricated metal products). Pre-treatment period: 2016–2018, before OSHA’s crystalline silica standard’s engineering controls deadline took effect for general industry (June 2021, first full year 2022). Source: OSHA Injury Tracking Application Form 300A.

as (count \div total hours worked) \times 200,000. Rates are winsorized at the 1st and 99th percentiles to limit the influence of establishments with very few reported hours.

Table 1 reports pre-treatment (2016–2021) means. High-silica establishments are smaller (102 vs. 227 employees) but have comparable respiratory condition rates (0.017 vs. 0.018 per 100 FTE). High-silica establishments have notably higher hearing loss rates (0.215 vs. 0.109), consistent with the noisier production environments in metalworking and stone fabrication.

4. Empirical Strategy

The core test reshapes the data into a panel at the establishment-by-hazard-by-year level, stacking five hazard categories (total injuries, respiratory conditions, hearing loss, skin disorders, other illnesses) for each establishment-year. The triple-difference specification is:

$$Y_{icht} = \alpha_{ic} + \gamma_{ct} + \lambda_{it} + \beta_3 (\text{HighSilica}_i \times \text{NonTargeted}_c \times \text{Post}_t) + \varepsilon_{icht} \quad (1)$$

where i indexes establishments, c indexes hazard categories, and t indexes years. HighSilica_i equals one for NAICS 327, 331, or 332. NonTargeted_c equals one for all hazard categories except respiratory conditions. Post_t equals one for 2022 onward (first full year after the June 2021 engineering controls deadline). The saturated fixed effects absorb: establishment-by-hazard profiles (α_{ic}), national hazard-by-year trends (γ_{ct} , including COVID effects on

specific illness categories), and establishment-by-year shocks (λ_{it} , including firm-level COVID impacts and reporting changes).

The coefficient β_3 identifies cross-hazard substitution: a positive estimate would indicate that non-targeted hazard categories worsened differentially at high-silica establishments—relative to targeted hazards, relative to low-silica manufacturers, and relative to the pre-regulation period. Standard errors are clustered at the establishment level.

Identifying assumptions. The triple-difference requires that, absent the silica standard, the gap between targeted and non-targeted hazard categories would have evolved similarly at high-silica and low-silica manufacturers. This is weaker than a standard DiD parallel-trends assumption because the establishment-by-year fixed effects absorb any establishment-specific time shock that affects all hazard categories equally—including COVID, reporting-rule changes, and business-cycle variation. Only a shock that differentially affects non-respiratory versus respiratory conditions at high-silica versus low-silica establishments, timed with the regulation, would confound the estimate.

I also estimate individual DiD specifications for each hazard category:

$$Y_{it} = \alpha_i + \delta_t + \beta (\text{HighSilica}_i \times \text{Post}_t) + \varepsilon_{it} \quad (2)$$

and an event study replacing Post_t with year indicators (base: 2018) interacted with $\text{HighSilica}_i \times \text{NonTargeted}_c$.

5. Results

5.1 Main Results

Table 2 presents the central findings. Column 1 reports the triple-difference estimate from Equation (1): $\hat{\beta}_3 = -0.032$ with a standard error of 0.010 ($p = 0.002$). The negative sign is the opposite of what cross-hazard substitution would predict. Non-targeted hazard categories improved *more* than the targeted category at high-silica establishments after the engineering controls deadline.

Columns 2 and 3 decompose the result by hazard category. Respiratory conditions declined modestly at high-silica establishments (-0.002 per 100 FTE, Column 2, $p = 0.11$), consistent with much of the respiratory improvement occurring during the PEL phase-in years (2018–2021). Total injury rates declined significantly: -0.135 per 100 FTE (Column 3, $p = 0.002$).

Critically, the negative triple-difference is driven entirely by total injuries. When I restrict

Table 2: The Safety Balloon: Does Targeted Regulation Create Cross-Hazard Substitution?

	(1) Triple-Diff All Hazards	(2) Respiratory (Targeted)	(3) Total Injuries (Non-Targeted)
High-silica \times Non-targeted \times Post	-0.0322** (0.0103)		
High-silica \times Post		-0.0018 (0.0011)	-0.1346** (0.0429)
Establishment FE	✓	✓	✓
Estab. \times Hazard FE	✓		
Hazard \times Year FE	✓		
Estab. \times Year FE	✓		
Year FE		✓	✓
Observations	1,274,265	254,853	254,853

Notes: Column 1 reports the triple-difference estimate: the differential change in non-targeted hazard categories (injuries, hearing loss, skin disorders, other illnesses) relative to the targeted category (respiratory conditions) at high-silica establishments vs. low-silica establishments, before vs. after the 2018 silica standard. Columns 2–3 report difference-in-differences for individual hazard categories. Rates per 100 FTE-years. Standard errors clustered at establishment level in parentheses. *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$.

the non-targeted category to illness types only (hearing loss, skin disorders, other illnesses), the triple-difference becomes +0.002 (SE = 0.003, $p = 0.51$)—a precise zero. Category-specific DiDs confirm: hearing loss (-0.005 , $p = 0.28$), skin disorders (-0.002 , $p = 0.06$), and other illnesses ($+0.006$, $p = 0.15$) all show small, insignificant changes. The safety balloon does not inflate for any illness category. The complementarity with total injuries likely reflects engineering controls (enclosed workstations, ventilation barriers) that reduce both silica exposure and physical injury risk simultaneously.

5.2 Event Study

Table 3 reports the event-study decomposition. At $t - 5$ (2016, the first ITA reporting year), the coefficient is near zero (-0.005 , $p = 0.78$), consistent with no differential pre-existing trend. During the PEL phase-in period, coefficients become increasingly negative ($t - 4$ through $t - 2$), reflecting the gradual onset of compliance investments. The immediate pre-treatment coefficient ($t - 1$, covering 2020) is -0.013 (SE = 0.017), insignificant.

After the engineering controls deadline, the coefficients grow strongly negative: +0.025 at $t + 1$ (2022, insignificant), -0.090 at $t + 2$ (2023, $p < 0.001$), and -0.126 at $t + 3$ (2024, $p < 0.001$). This accelerating pattern is consistent with cumulative safety infrastructure

Table 3: Event Study: Triple-Difference Coefficients by Year

Event Time (Year)	Estimate	SE
$t - 5$ (2016)	-0.0053	(0.0193)
$t - 4$ (2017)	-0.0435*	(0.0183)
$t - 3$ (2018)	-0.0431*	(0.0176)
$t - 2$ (2019)	-0.0523**	(0.0164)
$t - 1$ (2020)	-0.0131	(0.0167)
$t + 1$ (2022)	0.0254	(0.0163)
$t + 2$ (2023)	-0.0905***	(0.0172)
$t + 3$ (2024)	-0.1262***	(0.0182)
$t = 0$ (2021)	[Base year]	
Observations	1,274,265	
Estab. \times Hazard + Hazard \times Year + Estab. \times Year FE	✓	

Notes: Each coefficient is the triple-difference interaction (High-silica \times Non-targeted hazard) at event time k relative to $t = 0$ (2018). Pre-treatment years $t - 2$ and $t - 1$ test for differential pre-trends. COVID pandemic in 2020–2021 may affect coefficients at $t + 2$ and $t + 3$. Standard errors clustered at establishment level. *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$.

improving multiple hazard dimensions simultaneously, with the full effect materializing as engineering controls are fully installed and operational.

5.3 Robustness

Table 4 shows the triple-difference is robust to alternative treatment definitions but sensitive to the sample period. Broadening the silica definition to include NAICS 321 (wood) and 324 (petroleum) strengthens the result to -0.052 ($p < 0.001$). Narrowing to NAICS 327 alone yields -0.049 ($p = 0.015$). In no specification does the triple-difference turn positive—the safety balloon never inflates.

However, *excluding COVID years 2020–2021 attenuates the coefficient from -0.032 to -0.003* (Panel A, Row 2). This is an important finding: the total-injury complementarity appears concentrated in the pandemic era, when many manufacturers installed ventilation, PPE programs, and health monitoring systems that overlapped with silica compliance infrastructure. The core null finding for cross-hazard substitution *among illness categories* is unaffected by this concern, as it holds regardless of COVID inclusion. But the stronger claim of complementarity through total injuries should be treated with caution.

Panel B reports alternative outcome measures. Days away from work declined by 0.82 per 100 FTE at high-silica establishments ($p = 0.39$). Log injury counts fell by 0.9% ($p = 0.16$). These are directionally consistent but imprecise.

Table 4: Robustness of the Cross-Hazard Substitution Test

Specification	Estimate	SE	<i>N</i>
<i>Panel A: Triple-Difference Variants</i>			
Baseline (Table 2, Col. 1)	-0.0322	(0.0103)	1,274,265
Excl. COVID years (2020–2021)	-0.0032	(0.0128)	923,185
Broader silica def. (add NAICS 321/324)	-0.0516	(0.0098)	1,274,265
Narrower silica def. (NAICS 327 only)	-0.049	(0.020)	1,274,265
<i>Panel B: Alternative Outcome Measures</i>			
Days away from work (DAFW) rate	-0.820	(0.956)	254,853
Log(injuries + 1)	-0.0089	(0.0064)	254,853

Notes: Panel A varies the triple-difference specification. All include establishment \times hazard, hazard \times year, and establishment \times year fixed effects. Panel B reports DiD for high-silica \times post with establishment and year FE only. Standard errors clustered at establishment level. *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$.

5.4 Heterogeneity by Firm Size

If cross-hazard substitution arises from binding budget constraints, it should be stronger at smaller firms with less financial slack. I split the triple-difference by establishment size (above and below 500 employees). The small-firm coefficient is -0.034 ($p = 0.001$)—small firms show *stronger* complementarity, not substitution. The large-firm coefficient is $+0.048$ (SE = 0.029, $p = 0.10$), borderline positive but insignificant. The pattern—complementarity concentrated at small firms—is inconsistent with the budget-constraint mechanism: the firms predicted to be most constrained are the ones showing the strongest spillovers.

6. Discussion

Why didn't the balloon inflate? Three mechanisms could explain the complementarity.

First, *positive spillovers from compliance infrastructure*. Engineering controls installed for silica—ventilation systems, enclosed cutting stations, wet suppression—simultaneously reduce other hazards. Improved dust collection reduces both silica exposure and general particulate levels; enclosed workstations reduce both respiratory and injury risk from flying debris. The phased event study is consistent with this channel: the complementarity strengthens as capital-intensive engineering controls come online.

Second, *managerial attention complementarities*. The compliance process—exposure assessments, training, medical surveillance—may heighten general safety awareness. [Johnson et al. \(2023\)](#) find that OSHA inspections reduce injuries beyond the specific violations cited,

suggesting a “general deterrence” channel that operates through management attention rather than specific abatement. The concentration of complementarity at small firms (where a single safety manager covers all hazards) supports this interpretation.

Third, *safety culture effects*. Firms that invest in one dimension of safety may develop institutional norms that promote safety broadly. The silica standard required firms to designate competent persons, create written exposure control plans, and train workers—organizational changes that may generalize beyond the specific hazard.

The attenuation when COVID years are excluded (from -0.032 to -0.003) deserves attention. The COVID pandemic forced many manufacturers to install ventilation, PPE programs, and health monitoring systems—investments that overlap substantially with silica compliance requirements. The complementarity may partly reflect COVID-era safety investments piggy-backing on silica compliance infrastructure, a channel that would not replicate in non-pandemic settings.

These results carry a direct policy implication. OSHA’s shift toward hazard-specific standards—silica, beryllium, heat—does not appear to create the cross-hazard tradeoffs that critics feared. Targeted regulation produced broad-based safety improvements, consistent with the “rising tide” view of safety culture rather than the “fixed pie” view of safety budgets.

7. Conclusion

Squeeze one end of the balloon, and the other end is supposed to inflate. In the case of OSHA’s silica standard, it deflated everywhere. High-silica manufacturers that invested in respiratory protections saw their non-targeted injury categories improve *more* than the targeted one—the opposite of the substitution feared by critics of hazard-specific regulation. The safety balloon hypothesis finds no support in 1.3 million establishment-hazard-year observations spanning nine years of American manufacturing. Whether the complementarity arises from shared infrastructure, managerial attention, or safety culture remains an open question for future work with richer firm-level data on compliance investments. But the first-order finding is clear: for the canonical case of a major U.S. occupational health standard, targeted regulation did not merely avoid tradeoffs—it generated spillovers that improved workplace safety broadly.

Acknowledgements

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

Project Repository: <https://github.com/SocialCatalystLab/ape-papers>

Contributors: @ai1scl

First Contributor: <https://github.com/ai1scl>

References

- Bartel, Ann P and Lacy Glenn Thomas**, “OSHA’s Impact on Safety,” *Journal of Law and Economics*, 1985, 28 (1), 29–63.
- Gray, Wayne B and John T Scholz**, “The Declining Effects of OSHA Inspections on Manufacturing Injuries, 1979–1998,” *ILR Review*, 2005, 58 (4), 571–587.
- Johnson, Matthew S**, “Regulation by Shaming: Deterrence Effects of Publicizing Violations of Workplace Safety and Health Laws,” *American Economic Review*, 2020, 110 (6), 1866–1904.
- , **David I Levine, and Michael W Toffel**, “Do Investments in Safety Lead to Fewer Worker Injuries? Evidence from OSHA Inspections,” *American Economic Journal: Applied Economics*, 2023, 15 (3), 29–63.
- Levine, David I, Michael W Toffel, and Matthew S Johnson**, “Randomized Government Safety Inspections Reduce Worker Injuries with No Detectable Job Loss,” *Science*, 2012, 336 (6083), 907–911.
- OSHA**, “Occupational Exposure to Respirable Crystalline Silica; Final Rule,” Federal Register 81 FR 16286, Occupational Safety and Health Administration 2016.
- Peltzman, Sam**, “The Effects of Automobile Safety Regulation,” *Journal of Political Economy*, 1975, 83 (4), 677–725.
- Shimshack, Jay P**, “The Economics of Environmental Monitoring and Enforcement,” *Annual Review of Resource Economics*, 2014, 6, 339–360.
- Viscusi, W Kip**, “The Impact of Occupational Safety and Health Regulation,” *The Bell Journal of Economics*, 1979, 10 (1), 117–140.
- , “The Lulling Effect: The Impact of Child-Resistant Packaging on Aspirin and Analgesic Ingestions,” *American Economic Review*, 1984, 74 (2), 324–327.
- , *Fatal Tradeoffs: Public and Private Responsibilities for Risk*, Oxford University Press, 1992.

A. Data Appendix

The analysis uses OSHA’s Injury Tracking Application (ITA) Form 300A summary data, downloaded from OSHA’s establishment-specific data portal¹ for reporting years 2016–2024. Each record represents one establishment-year and includes: establishment identifiers (name, address, NAICS code, EIN), employment (annual average employees, total hours worked), and injury/illness counts disaggregated into six categories (total injuries, skin disorders, respiratory conditions, poisonings, hearing loss, other illnesses) plus severity measures (days away from work, days of job transfer/restriction, deaths).

Sample construction. Starting from 2,803,410 raw establishment-year observations, I apply the following filters:

1. Drop observations with missing or zero annual average employees or total hours worked: 2,792,308 remain.
2. Restrict to manufacturing (NAICS 2-digit codes 31–33): 893,145 observations from 141,682 establishments.
3. Require establishment presence in ≥ 4 of 9 reporting years: 254,933 observations from 43,954 establishments.

Treatment classification. High-silica subsectors are defined as NAICS 327 (nonmetallic mineral products: 4,122 establishments), NAICS 331 (primary metals: 3,891), and NAICS 332 (fabricated metal products: 10,458). These three subsectors account for the majority of silica exposure in general industry per OSHA’s regulatory impact analysis. All other manufacturing subsectors (NAICS 311–326, 333–339) form the low-silica comparison group.

Rate construction. Injury and illness rates are expressed per 100 full-time equivalent (FTE) workers:

$$\text{Rate} = \frac{\text{Count}}{\text{Total Hours Worked}} \times 200,000$$

where 200,000 represents 100 workers \times 2,000 annual hours. Rates are winsorized at the 1st and 99th percentiles of the manufacturing distribution.

B. Robustness Appendix

Additional robustness checks confirm the main findings. The construction sector—subject to its own silica standard effective June 2016—serves as a separate placebo test: a triple-difference

¹<https://www.osha.gov/Establishment-Specific-Injury-and-Illness-Data>

Table 5: Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
<i>Panel A: Pooled</i>						
Cross-hazard substitution (DDD)	-0.0322	0.0103	2.326	-0.0138	0.0044	Small negative
Respiratory conditions (DiD)	-0.0018	0.0011	0.034	-0.0531	0.0328	Moderate negative
Total injuries (DiD)	-0.1346	0.0429	4.682	-0.0287	0.0092	Small negative
Days away from work (DiD)	-0.8201	0.9564	84.056	-0.0098	0.0114	Small negative
<i>Panel B: Heterogeneous (firm size, cross-hazard substitution)</i>						
Small firms (<500 emp.)	-0.0338	0.0104	2.326	-0.0145	0.0045	Small negative
Large firms (\geq 500 emp.)	0.0481	0.0294	2.326	0.0207	0.0126	Small positive

Notes: **Country:** United States. **Research question:** Does OSHA’s 2018 crystalline silica standard, which targeted respiratory hazards in manufacturing, cause cross-hazard injury substitution to non-targeted hazard categories within the same establishments? **Policy mechanism:** The standard reduced the permissible exposure limit for respirable crystalline silica from 250 to 50 $\mu\text{g}/\text{m}^3$ in general industry, requiring engineering controls, respiratory protection programs, medical surveillance, and exposure assessments—imposing substantial compliance costs on silica-intensive manufacturers. **Outcome definition:** Injury/illness rates per 100 full-time equivalent workers (count \div total hours worked \times 200,000), separately by hazard category (total injuries, respiratory conditions, hearing loss, skin disorders, other illnesses). **Treatment:** Binary: high-silica manufacturing (NAICS 327, 331, 332) versus low-silica manufacturing (all other NAICS 31–33 subsectors). **Data:** OSHA Injury Tracking Application Form 300A, 2016–2024, establishment-year level, 254,933 observations from 43,954 manufacturing establishments. **Method:** Triple-difference (high-silica \times non-targeted hazard \times post-2021) with establishment \times hazard, hazard \times year, and establishment \times year fixed effects; standard errors clustered at establishment level. **Sample:** Manufacturing establishments (NAICS 31–33) present in the ITA for \geq 4 years with positive employment and hours; excludes construction and services. $\text{SDE} = \hat{\beta}/\text{SD}(Y)$ where $\text{SD}(Y)$ is the pre-treatment standard deviation. Classification refers to magnitude, not statistical significance: Large ($|\text{SDE}| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).

comparing construction versus services shows a large positive coefficient (+0.470), but this reflects differential COVID impacts on healthcare workers (NAICS 62) in the service-sector control group rather than cross-hazard substitution. When restricted to manufacturing-only comparisons, all substitution estimates remain economically small and statistically insignificant.

C. Standardized Effect Sizes