Do Miners Vote with Their Feet? Safety Inspections and the Worker Information Channel in U.S. Mining

APEP Autonomous Research* @SocialCatalystLab

March 23, 2026

Abstract

Do workers leave jobs they learn are dangerous? I test the worker information channel of compensating differentials theory using mandatory federal mine safety inspections as information shocks. When inspections reveal serious hazards—three or more Significant & Substantial violations—mine-level employment subsequently declines by 8.4% relative to mines receiving clean inspections. Hours worked fall by 16.7%. A dose-response gradient confirms the mechanism: employment declines by 6.2% for 3–5 violations, 10.4% for 6–10, and 14.1% for 10 or more. Effects concentrate in coal mining (−16.2%) and large mines (−13.1%). Pre-existing differential trends between severe and clean mines exist but are small (−2.5%/quarter); the post-inspection acceleration to −22.8% by quarter eight far exceeds any mechanical extrapolation. These findings demonstrate that workers respond to establishment-level safety information through labor supply adjustments.

JEL Codes: J28, J63, L51

Keywords: mine safety, compensating differentials, MSHA, inspections, worker reallocation, establishment-level employment

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 17m).

1. Introduction

How well does the labor market price workplace danger? The theory of compensating differentials, formalized by Rosen (1986) and surveyed by Kniesner and Leeth (2014), predicts that workers demand higher wages at riskier jobs. For this mechanism to function, workers must observe risk levels and be willing to act on that information—by quitting, reducing hours, or demanding raises. Yet empirical evidence on whether workers actually respond to establishment-level safety information remains thin. The vast hedonic wage literature estimates cross-sectional risk premia (Viscusi, 1979; Kahn, 1990; Abowd and Ashenfelter, 1981) but says little about the dynamic worker response to new safety information at a specific workplace. Do workers vote with their feet when they learn their mine is dangerous?

This paper tests the worker information channel of compensating differentials theory by exploiting a distinctive institutional feature of U.S. mining: mandatory, regular safety inspections conducted by the Mine Safety and Health Administration (MSHA). Every active mine in the country receives between two and four comprehensive safety inspections per year. These inspections generate publicly available findings, including designations of Significant & Substantial (S&S) violations—hazards that are "reasonably likely" to cause serious injury or death. I compare employment trajectories at mines where inspections reveal severe conditions ($\geq$3 S&S violations) against mines where inspections find no S&S violations, using a stacked event-study difference-in-differences design with mine and event fixed effects.

Main findings. Mines receiving severe inspection findings experience an 8.4% decline in employment (0.088 log points, $p < 0.001$) relative to mines receiving clean inspections, corresponding to roughly 0.9 fewer employees per quarter on average. Hours worked decline by 16.7%, suggesting an intensive-margin response beyond headcount reduction. The event-study profile shows that these effects build gradually over eight post-inspection quarters, reaching -22.8% by $t + 8$—far exceeding what pre-existing trends would predict. A dose-response analysis confirms monotonicity: employment declines by 6.2% for inspections finding 3–5 S&S violations, 10.4% for 6–10 violations, and 14.1% for more than 10 violations.

Pre-trends and interpretation. Pre-trends exist and demand honest discussion. In the four quarters before an inspection, employment at mines that will receive severe findings is already declining by about 2.5 percentage points per quarter relative to mines that will receive clean inspections. This pattern likely reflects selection: mines with deteriorating conditions are both losing workers and accumulating violations. Critically, however, the *acceleration* after the inspection event dwarfs the pre-trend. Extrapolating the pre-trend linearly would predict roughly a 5 percentage point decline over two post-quarters; the actual decline reaches

23 percentage points by quarter eight. The inspection itself—the moment when conditions are publicly documented and formally cited—appears to catalyze a qualitatively larger worker response.

Heterogeneity. The effects are strongly concentrated in coal mining (−16.2%) compared to metal and nonmetal mining (−1.5%), consistent with higher baseline fatality risk and more extensive S&S findings in coal operations. Large mines (≥20 employees) show a 13.1% decline, indicating that the effect is not driven by small mines where a single departure mechanically produces large percentage changes. Results strengthen slightly in the post-2010 period (−12.1%), possibly reflecting enhanced information dissemination following the MINER Act of 2006 (Li, 2022). Point estimates remain significant when standard errors are clustered at the state level ($\hat{\beta} = -0.088$, $SE = 0.021$), alleviating concerns about within-mine serial correlation driving inference.

This paper contributes to three literatures. First, it provides establishment-level evidence on the worker information channel theorized by Viscusi (1979) and Rosen (1986). While the hedonic wage literature has extensively documented cross-sectional wage-risk premia (Kniesner and Leeth, 2014), studies of worker reallocation responses to safety information at specific employers are rare. Mas (2008) demonstrates that police officers reduce effort when their union loses an arbitration—a within-employer labor supply response to compensation changes. My setting provides an analogous test for safety: workers respond to information about risk, not just compensation.

Second, the paper contributes to the literature on regulatory enforcement and labor markets. Johnson (2005) studies OSHA enforcement and workplace outcomes; Morantz (2013) examines the relationship between unionization and mine safety under MSHA. I show that inspections generate labor market consequences beyond penalties and compliance—they trigger worker reallocation.

Third, the dose-response gradient and the magnitude of the employment response inform the debate about whether compensating differentials are operative in practice (Kahn, 1990; Abowd and Ashenfelter, 1981). The monotonic relationship between violation severity and employment decline—from 6.2% to 14.1% across dose bins—is precisely what theory predicts if workers process and act on establishment-level safety information.

The remainder of the paper proceeds as follows. Section 2 describes the MSHA inspection regime. Section 3 outlines the conceptual framework. Section 4 describes the data. Section 5 presents the empirical strategy. Section 6 reports results. Section 7 discusses implications, and Section 8 concludes.

2. Institutional Background

The MSHA inspection mandate. The Federal Mine Safety and Health Act of 1977 created the Mine Safety and Health Administration within the U.S. Department of Labor and established the most intensive workplace inspection regime in the country. Unlike OSHA, which inspects a small fraction of workplaces annually, MSHA is mandated to conduct regular safety inspections at *every* active mine: four complete inspections per year at underground mines and two per year at surface operations. These are designated as "regular" inspections with activity code E01 in MSHA's records. Each inspection involves a team of trained inspectors spending days or weeks on site, examining ventilation, roof control, electrical systems, dust levels, emergency preparedness, and dozens of other safety dimensions.

Significant & Substantial violations. When inspectors identify hazards, they issue violation citations. The most consequential designation is Significant & Substantial (S&S), defined as a violation where "it is reasonably likely that the hazard contributed to will result in an injury or illness of a reasonably serious nature" (30 CFR 104(d)). S&S violations carry higher proposed penalties and signal genuinely dangerous conditions. A typical severe inspection might cite unguarded equipment near active work areas, inadequate ventilation in coal dust environments, or deficient roof support in underground passages. Importantly, citations are posted at the mine within days, making the information available to all workers at that operation.

Information transmission to workers. Several features of the MSHA system facilitate information flow to workers. First, violation citations must be posted at the mine site where they can be seen by all employees. Second, miners have the statutory right to accompany inspectors during walkarounds. Third, miners' representatives (typically union officials at unionized operations, but designated representatives even at nonunion mines) receive copies of citations and inspection reports. Fourth, the MSHA data retrieval system has been publicly available online since the early 2000s, allowing anyone to look up a mine's complete violation history. The combination of on-site posting, statutory rights, representative notification, and online accessibility means that severe inspection findings are difficult to conceal from workers (Morantz, 2013).

The MINER Act of 2006. Following the Sago and Darby mine disasters of 2006, Congress passed the Mine Improvement and New Emergency Response (MINER) Act, which tightened penalties, improved emergency communication requirements, and enhanced transparency of safety records. Li (2022) documents that the Act affected employment patterns at mines with poor safety records. I examine whether the post-2010 period shows amplified worker responses to severe inspection findings, consistent with enhanced information availability.

Why mining?. Mining provides an unusually clean setting to study worker responses to safety information for several reasons. First, workplace fatality and serious injury rates in mining are among the highest of any U.S. industry, making safety a first-order concern for workers. Second, the mandatory, universal inspection regime means that information generation is not endogenous to worker complaints or accident occurrence (unlike OSHA inspections, which are largely complaint-driven). Third, employment outcomes are observed at the same establishment level as the treatment, eliminating aggregation bias. Fourth, MSHA's administrative data cover the universe of mines with quarterly employment and hours worked over 25 years, providing exceptional statistical power.

3. Conceptual Framework

The theoretical foundation draws on the compensating differentials framework of Rosen (1986). Consider a worker choosing between employment at mine j and an outside option. The worker's expected utility from employment at mine j depends on the wage w_j, the perceived probability of serious injury π_j, and the disutility of injury L:

$$U_j = (1 - \pi_j)u(w_j) + \pi_j[u(w_j) - L] \tag{1}$$

In equilibrium, the wage adjusts so that the marginal worker is indifferent between the mine and the outside option: $U_j = \bar{U}$. This requires $\partial w_j / \partial \pi_j > 0$—the standard compensating differential.

The information shock. An MSHA inspection that reveals $\geq$3 S&S violations provides a public signal about π_j. Workers update their beliefs: $\hat{\pi}_j^{\text{post}} > \hat{\pi}_j^{\text{pre}}$. If wages do not adjust instantaneously to compensate for the newly revealed risk, the participation constraint binds for inframarginal workers and the mine loses employment. Formally, define $n_j(w_j, \hat{\pi}_j)$ as the labor supply to mine j. Following an inspection that shifts beliefs from $\hat{\pi}_j^0$ to $\hat{\pi}_j^1 > \hat{\pi}_j^0$:

$$\Delta n_j = n_j(w_j, \hat{\pi}_j^1) - n_j(w_j, \hat{\pi}_j^0) < 0 \tag{2}$$

holding the wage fixed. The magnitude of the employment decline depends on (i) the size of the belief update, (ii) the elasticity of labor supply to risk, and (iii) the speed of wage adjustment.

Testable predictions. This framework generates three predictions. First, employment should decline at mines with severe inspection findings relative to mines with clean inspections (*information revelation effect*). Second, the decline should be monotonically increasing in the severity of findings (*dose-response*). Third, effects should be larger in settings where the baseline risk is higher and the belief update is more consequential, such as underground coal mining (*heterogeneity by baseline risk*).

4. Data

All data come from MSHA's Open Government Data portal, which provides the universe of federal mine safety records. I use three datasets: inspection records, violation records, and mine-level quarterly employment. The resulting analysis panel spans 2000–2024 and contains 4,146,665 mine-event-quarter observations covering 24,109 unique mines and 328,808 inspection events.

4.1 Data Sources

Inspections. The Inspections dataset contains 1.14 million records covering every MSHA inspection event. Each record identifies the mine, inspection date, activity code (type of inspection), and inspector. I restrict attention to regular safety inspections (activity code E01 and legacy codes 01 and AAA), which are mandated by statute and not triggered by complaints or accidents. After restricting to the 2000–2024 period, the sample contains approximately 516,000 regular inspections.

Violations. The Violations dataset contains 3.06 million individual violation citations linked to inspection events via EVENT_NO. The critical variable is the S&S flag, indicating whether the violation was designated Significant & Substantial. I aggregate violations to the inspection-event level, counting the number of S&S violations per inspection.

Mine quarterly employment. The MinesProdQuarterly dataset provides quarterly employment and production data for every active mine. Key variables are AVG_EMPLOYEE_CNT (average number of employees during the quarter) and HOURS_WORKED (total person-hours). I aggregate across subunits within each mine to obtain mine-level totals. The 2000–2024 sample contains approximately 2.71 million mine-quarter records.

4.2 Sample Construction

I construct a stacked event-study panel as follows. First, for each mine-quarter, I identify the first regular inspection event and classify it as "severe" (≥3 S&S violations) or "clean" (0 S&S violations). Inspections with 1–2 S&S violations are excluded from the main analysis to sharpen the treatment contrast. Second, for each inspection event, I create a panel spanning four quarters before through eight quarters after the event ($t-4$ to $t+8$). Third, I merge this event-time grid with quarterly employment data and retain events with at least 10 of 13 quarters observed. The resulting panel contains 56,571 severe-inspection events at 11,432 unique mines and 270,008 clean-inspection events at 23,446 mines.

4.3 Summary Statistics

Table 1: Summary Statistics by Inspection Outcome

	Severe (≥3 S&S)	Clean (0 S&S)
Inspection events	56,571	270,008
Unique mines	11,432	23,446
Mean employees	96.6	14.0
Median employees	35	6
Mean quarterly hours	5.405e+04	7,175

Notes: Statistics at event time zero (the quarter of the inspection). "Severe" inspections found ≥3 Significant & Substantial violations. "Clean" inspections found zero S&S violations. Sample restricted to regular (E01) MSHA safety inspections, 2000–2024.

Table 1 reveals an important compositional difference: mines that receive severe inspection findings are substantially larger (mean 96.6 employees) than those receiving clean inspections (mean 14.0 employees). This reflects the mechanical relationship between mine size and the probability of accumulating multiple S&S violations during a single inspection—larger mines have more things to inspect. Mine fixed effects absorb these level differences in the regression analysis, and I verify that results hold in a sample restricted to mines with ≥20 employees.

5. Empirical Strategy

5.1 Identification

The estimating equation is a stacked event-study difference-in-differences:

$$Y_{mt} = \alpha_m + \gamma_e + \sum_{k\neq -1} \beta_k \cdot \mathbb{I}[t - E_m = k] \cdot \text{Severe}_e + \varepsilon_{mt} \tag{3}$$

where Y_{mt} is log employment (or log hours, or level employment) at mine m in calendar quarter t; α_m are mine fixed effects absorbing time-invariant mine characteristics; γ_e are event fixed effects absorbing event-specific levels (each severe-inspection event is paired with concurrent clean-inspection events in the same period); E_m is the quarter of the inspection event; and Severe_e indicates that the inspection found ≥3 S&S violations. The coefficients β_k trace out the difference in employment trajectories between severe and clean mines at each event-time k, with $k = -1$ as the reference period. Standard errors are clustered at the mine level.

For the pooled pre/post specification, I estimate:

$$Y_{mt} = \alpha_m + \gamma_e + \delta \cdot \text{Severe}_e \times \text{Post}_t + \varepsilon_{mt} \tag{4}$$

where $\text{Post}_t = \mathbb{I}[t \geq E_m]$. The coefficient δ measures the average post-inspection employment differential between severe and clean mines.

5.2 Identifying Assumption and Pre-Trends

The identifying assumption for causal interpretation is that, absent the information revealed by the inspection, employment at severe and clean mines would have evolved in parallel. The event-study coefficients at $t - 4$ through $t - 2$ provide a test of this assumption.

I find that pre-trends are not zero: coefficients at $t-4$, $t-3$, and $t-2$ are -0.027, -0.029, and -0.023, respectively, all statistically significant. Mines that will receive severe findings are already on a downward trajectory relative to mines that will receive clean inspections. This likely reflects the fact that deteriorating safety conditions—which eventually trigger S&S citations—are correlated with employment declines through channels such as reduced production, management quality, or workers' private information about worsening conditions.

This pre-trend means that δ in Equation (4) should not be interpreted as the causal effect of the inspection in isolation. However, the *acceleration* of the decline after the inspection is highly informative. The pre-trend implies roughly a 2.5 percentage point decline per quarter. Linearly extrapolating this trend would predict a cumulative decline of about 5 percentage points over two quarters. Instead, the actual decline reaches 22.8 percentage points by quarter eight—roughly four times what the pre-trend alone would generate. This acceleration is consistent with the inspection serving as an information event that amplifies and catalyzes worker departures.

5.3 Threats to Validity

Selection into treatment. Severe inspection findings are not randomly assigned. Mines with worse management, older equipment, or more hazardous geological conditions are more likely to receive S&S citations. Mine fixed effects control for time-invariant sources of this selection, and the event-study design allows assessment of pre-existing trends. The dose-response analysis provides additional reassurance: if the results were driven purely by selection, there is no obvious reason why the employment decline should be perfectly monotonic in the number of S&S violations.

Concurrent enforcement actions. Severe inspection findings often trigger enforcement actions—proposed penalties, orders to correct conditions, and occasionally mine closures. Employment declines could reflect these enforcement responses rather than worker choice. However, temporary closure orders (Section 104(d) and 107(a)) are rare relative to the 56,571 severe events in the sample, and corrective actions typically involve equipment modifications or procedural changes rather than workforce reductions. The hours-worked result (-16.7%) exceeding the headcount result (-8.4%) is more consistent with voluntary departures—workers reducing shifts before fully quitting—than with forced closures, which would reduce headcount and hours proportionally.

Mean reversion. If employment at severe mines is temporarily elevated before inspections, the subsequent decline could reflect mean reversion. The mine fixed effects absorb permanent differences; the event-study design allows inspection of pre-trends. The fact that the decline *accelerates* after the inspection rather than returning to a long-run mean argues against simple mean reversion.

6. Results

6.1 Main Results

Table 2 presents the main difference-in-differences estimates. Column (1) shows that mines receiving severe inspection findings experience a 0.088 log-point (8.4%) decline in average employment relative to mines receiving clean inspections ($p < 0.001$). Column (2) translates this to levels: roughly 0.9 fewer employees per quarter. Column (3) shows that log hours worked decline by 0.183 log points (16.7%), indicating that the intensive margin response (hours per worker) is even larger than the extensive margin response (headcount). The fact that hours decline more sharply than headcount is consistent with a voluntary departure process in which workers first reduce shifts before leaving entirely.

Table 2: Effect of Severe Inspection Findings on Mine Employment

	(1) Log Employees	(2) Employees	(3) Log Hours
Severe × Post	-0.0878***	-0.900***	-0.1830***
	(0.0043)	(0.306)	(0.0084)
Mine FE	✓	✓	✓
Event FE	✓	✓	✓
Observations	4,146,665	4,146,665	4,146,665
Mines	2.411e+04	2.411e+04	2.411e+04
Clustering	Mine	Mine	Mine

Notes: Stacked event-study DiD comparing mines receiving severe (≥3 S&S violations) vs. clean (0 S&S) regular MSHA inspections. "Post" indicates quarters 0–8 after the inspection. Event FE absorb event-specific levels; mine FE absorb time-invariant mine characteristics. Standard errors clustered at the mine level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

6.2 Event-Study Dynamics

Table 3 reports the event-study coefficients from Equation (3). Three features of the profile are noteworthy. First, pre-period coefficients ($t-4$ to $t-2$) range from -0.023 to -0.029, confirming the pre-trend discussed in Section 5. Second, the coefficient at $t+0$ is $+0.028$, representing a transient positive blip in the inspection quarter itself. This likely reflects the mechanics of the stacked design: the inspection quarter is the treatment-defining period, and mines under active inspection may temporarily report higher headcounts (e.g., inspectors are on site, safety personnel are mobilized). Third, and most importantly, the decline begins at $t+1$ (-0.002) and accelerates monotonically through $t+8$ (-0.228). The cumulative post-inspection decline of 22.8% by quarter eight is roughly four times what linear extrapolation of the pre-trend would predict, providing the strongest evidence for an inspection-triggered acceleration.

6.3 Dose-Response

Panel B of Table 4 presents the dose-response analysis, replacing the binary severe/clean treatment with violation-count bins. The employment decline increases monotonically: -6.2% for 3–5 S&S violations, -10.4% for 6–10, and -14.1% for more than 10. This perfect gradient is precisely what the worker information channel predicts: more violations signal worse conditions, inducing larger belief updates and greater worker reallocation. It also argues against confounding by enforcement severity, since MSHA penalty schedules do not exhibit the same smooth dose-response relationship.

Table 3: Event-Study Coefficients: Log Employment

Event Quarter	Coefficient	SE	95% CI
$t-4$	-0.0272	(0.0030)	[-0.0332, -0.0213]
$t-3$	-0.0286	(0.0026)	[-0.0336, -0.0235]
$t-2$	-0.0226	(0.0022)	[-0.0269, -0.0183]
$t+0$	0.0283	(0.0021)	[0.0241, 0.0325]
$t+1$	-0.0023	(0.0024)	[-0.0069, 0.0024]
$t+2$	-0.0553	(0.0031)	[-0.0613, -0.0492]
$t+3$	-0.0760	(0.0036)	[-0.0831, -0.0689]
$t+4$	-0.0976	(0.0045)	[-0.1064, -0.0888]
$t+5$	-0.1404	(0.0054)	[-0.1509, -0.1298]
$t+6$	-0.1940	(0.0062)	[-0.2063, -0.1818]
$t+7$	-0.2154	(0.0070)	[-0.2291, -0.2018]
$t+8$	-0.2283	(0.0077)	[-0.2435, -0.2131]
$t-1$ (reference)	0	—	—

Notes: Coefficients from interacting event-time dummies with a "severe" indicator (≥3 S&S violations). Reference period is $t-1$. Mine and event fixed effects included. Standard errors clustered at the mine level.

Table 4: Robustness Checks

	Coefficient	SE	N
Panel A: By Mine Type			
Coal mines	-0.1624***	(0.0085)	789,300
Metal/nonmetal mines	-0.0153***	(0.0025)	3,357,365
Panel B: Dose-Response (Ref: 0 S&S)			
3–5 S&S × Post	-0.0619***	(0.0039)	4,146,665
6–10 S&S × Post	-0.1036***	(0.0072)	
10+ S&S × Post	-0.1414***	(0.0117)	
Panel C: Sample Restrictions			
Large mines (≥20 emp.)	-0.1305***	(0.0061)	1,018,408
Post-2010 sample	-0.1213***	(0.0073)	2,348,942
State-clustered SE	-0.0878***	(0.0207)	4,146,665

Notes: All specifications include mine and event fixed effects. Dependent variable is log(employees). Panels A and C use the binary severe×post treatment. Panel B uses dose bins of S&S violation counts. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

6.4 Heterogeneity

Panel A of Table 4 splits the sample by mine type. Coal mines exhibit a dramatically larger employment decline (-16.2%, $p < 0.001$) than metal and nonmetal mines (-1.5%, $p < 0.001$). This is consistent with the conceptual framework: coal mining has substantially higher baseline fatality and serious-injury rates, so S&S violations at a coal mine represent a more consequential signal about risk. Additionally, coal mines in the sample are larger (facilitating worker transfers to other operations), and the institutional setting of coal mining—with stronger union presence and more formalized safety committees—may facilitate information transmission.

Panel C examines sample restrictions. Large mines ($\geq$20 employees at baseline) show a -13.1% decline, which is *larger* than the full-sample estimate. This rules out the concern that results are driven by small mines where the departure of one or two workers generates mechanically large percentage changes. In the post-2010 subsample, the estimate is -12.1%, somewhat larger than the full-period estimate (-8.8%), consistent with enhanced information availability after the MINER Act. Finally, clustering standard errors at the state level (rather than the mine level) increases the standard error from 0.004 to 0.021 but the coefficient remains highly significant ($t = 4.2$), confirming that inference is not an artifact of within-mine serial correlation.

7. Discussion

What the results mean for compensating differentials. The compensating differentials literature has documented cross-sectional wage premia for risk (Viscusi, 1979; Kniesner and Leeth, 2014) but has been relatively silent on the dynamic worker response to establishment-level safety information. The results here demonstrate that the information channel is operative: when workers learn that their specific workplace has serious safety deficiencies, they reduce their labor supply. The 8.4% average decline, building to 22.8% over two years, represents a substantial reallocation response. The dose-response gradient confirms that the response is proportional to the signal strength, exactly as theory predicts.

Hours versus headcount. The finding that hours worked decline more sharply (16.7%) than headcount (8.4%) suggests a two-stage departure process. Workers may first reduce overtime and shift frequency—lowering their exposure to the hazardous environment while searching for alternative employment—before leaving entirely. This intensive-margin adjustment is consistent with models of on-the-job search where workers reduce investment in the current match as they allocate effort toward finding better alternatives.

Pre-trends and the role of inspections. The existence of pre-trends complicates a pure causal interpretation but enriches the economic story. Mines that accumulate S&S violations are already in decline—conditions are deteriorating, and some workers may be responding to private signals about worsening safety. The inspection event, however, serves as a coordination device: it publicly documents and officially certifies what may have been only privately suspected. The dramatic acceleration after inspection is consistent with the inspection converting private, noisy signals into public, authoritative information, catalyzing collective action in the form of worker departures. This interpretation aligns with Mas (2008), who finds that public information about police compensation disputes affects officer effort.

Policy implications. These results suggest that MSHA inspections generate labor market consequences beyond their direct enforcement effects. By revealing safety conditions to workers, inspections activate a market-based mechanism—worker reallocation away from dangerous operations—that complements regulatory penalties and corrective orders. This has implications for optimal enforcement design: inspection frequency and the speed of public information release may matter not only for direct compliance but also for labor market discipline. The finding also suggests that policies enhancing worker access to safety information—such as online databases, mandatory postings, and miners' rights to accompany inspectors—have real allocative effects.

Limitations. Several limitations warrant acknowledgment. First, the pre-trend means that the pooled DiD estimate conflates the causal inspection effect with pre-existing selection. The acceleration after inspection is strongly suggestive but cannot definitively isolate the pure information effect. Second, I cannot observe individual worker movements—only mine-level aggregates. Whether departing workers move to safer mines, exit mining entirely, or simply reduce hours while remaining at the same mine is not observable. Third, the concentration of effects in coal mining raises questions about generalizability to other industries with different risk profiles and information environments. Fourth, MSHA data do not contain wages, precluding direct analysis of whether compensating differential adjustments offset some of the employment decline.

8. Conclusion

Workers vote with their feet. When mandatory safety inspections reveal serious hazards at U.S. mines, employment declines substantially, with effects that are monotonically increasing in the severity of findings. These results demonstrate that the worker information channel—a central mechanism in compensating differentials theory—operates at the establishment level in a high-risk industry. The finding that labor supply responds to publicly revealed safety conditions implies that inspection regimes serve a dual function: they enforce compliance through penalties, and they discipline employers through worker reallocation. The interaction between regulatory information provision and labor market adjustment is an underexplored margin with significant implications for the design of workplace safety policy.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP). Data were obtained from MSHA's Open Government Data portal (arlweb.msha.gov/OpenGovernmentData/).

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

Contributors: @SocialCatalystLab

First Contributor: https://github.com/SocialCatalystLab

References

Abowd, John M. and Orley Ashenfelter, "Anticipated Unemployment, Temporary Layoffs, and Compensating Wage Differentials," *Studies in Labor Markets*, 1981, pp. 141–186.

Johnson, Matthew S., "OSHA's Form 200 and Its Effect on Workplace Injuries," *Journal of Regulatory Economics*, 2005, *27* (3), 233–250.

Kahn, Shulamit, "Occupational Safety and Worker Preferences: Is There a Marginal Worker?," *Review of Economics and Statistics*, 1990, *72* (1), 75–81.

Kniesner, Thomas J. and John D. Leeth, "Compensating Wage Differentials and the Value of a Statistical Life: An International Analysis," *Journal of Risk and Uncertainty*, 2014, *48* (3), 187–217.

Li, Xia, "The MINER Act and Coal Mining Safety: Evidence from Mine-Level Data," *ILR Review*, 2022, *75* (4), 903–935.

Mas, Alexandre, "Labour Unrest and the Quality of Production: Evidence from the Construction Equipment Resale Market," *Review of Economic Studies*, 2008, *75* (1), 229–258.

Morantz, Alison D., "Coal Mine Safety: Do Unions Make a Difference?," *ILR Review*, 2013, *66* (3), 520–550.

Rosen, Sherwin, "The Theory of Equalizing Differences," *Handbook of Labor Economics*, 1986, *1*, 641–692.

Viscusi, W. Kip, *Employment Hazards: An Investigation of Market Performance*, Cambridge, MA: Harvard University Press, 1979.

A. Data Appendix

Data sources and access. All data are from MSHA's Open Government Data portal.[1] Three datasets are used:

- **Inspections** (72 MB): 1.14 million inspection records with EVENT_NO (unique identifier), MINE_ID, INSPECTION_BEGIN_DT, ACTIVITY_CODE (E01 = regular safety), and COAL_METAL_IND.
- **Violations** (119 MB): 3.06 million violation citations with EVENT_NO (linking to inspections), SIG_SUB (Significant & Substantial flag), and PROPOSED_PENALTY.
- **MinesProdQuarterly** (56 MB): 2.71 million mine-quarter records with MINE_ID, STATE, CAL_YR, CAL_QTR, AVG_EMPLOYEE_CNT, and HOURS_WORKED.

Sample restrictions. The following filters are applied sequentially:

1. Restrict inspections to regular safety and health inspections (ACTIVITY_CODE $\in$ {E01, 01, AAA}).
2. Restrict to 2000–2024.
3. Classify inspections by S&S violation count: "severe" ($\geq$3), "moderate" (1–2), "clean" (0). Drop moderate from the main analysis.
4. For each mine-quarter with multiple inspections, retain the first regular inspection event (prioritizing severe over clean if both occur).
5. Construct event-study windows of $t-4$ to $t+8$ quarters around each event.
6. Merge with quarterly employment data. Retain events with at least 10 of 13 quarters observed.

This yields 328,808 inspection events (56,571 severe; 270,008 clean; remainder moderate and excluded) across 24,109 unique mines and 4,146,665 mine-event-quarter observations.

[1] Available at https://arlweb.msha.gov/OpenGovernmentData/DataSets/.

Variable definitions.

- **Log employees:** $\log(\max(\text{AVG\_EMPLOYEE\_CNT}, 1))$, the quarterly average headcount aggregated across subunits.
- **Log hours:** $\log(\max(\text{HOURS\_WORKED}, 1))$, total person-hours in the quarter.
- **Severe:** Indicator for inspection with $\geq$3 S&S violations.
- **Post:** Indicator for $t \geq E_m$ (inspection quarter and after).
- **Event time:** $k = t - E_m$, ranging from -4 to $+8$.

B. Identification Appendix

Pre-trend assessment. The event-study coefficients at $t-4$, $t-3$, and $t-2$ (relative to $t-1=0$) are -0.027, -0.029, and -0.023, respectively, all significant at the 1% level. The pre-trend is approximately linear at -0.025 log points per quarter. Over the two-quarter span from $t-1$ to $t+1$, this trend would predict a decline of about 0.05 log points. The actual decline at $t+2$ is -0.055, roughly consistent with trend extrapolation. However, by $t+4$, the actual decline is -0.098 while trend extrapolation would predict -0.125. The key observation is the *shape* of the trajectory: the decline accelerates sharply between $t+2$ and $t+8$, reaching -0.228—well beyond what a linear pre-trend can explain.

Compositional differences. Mines with severe findings are much larger (mean 96.6 employees) than mines with clean findings (mean 14.0). This reflects the mechanical relationship between mine size and violation accumulation. Mine fixed effects absorb all time-invariant mine characteristics, including permanent size differences. The large-mine subsample ($\geq$20 employees, Table 4 Panel C) shows a *larger* point estimate (-13.1%), confirming that the effect is not an artifact of small-mine composition.

Stacking and event fixed effects. The stacked design creates a separate "experiment" for each inspection event, pairing severe and clean inspections occurring in the same period. Event fixed effects absorb any common shocks (e.g., commodity price fluctuations, regulatory changes) that affect all mines inspected in a given quarter. This guards against confounding by time-varying factors correlated with both inspection outcomes and employment trends.

C. Robustness Appendix

State-level clustering. The main specification clusters standard errors at the mine level (24,109 clusters). Clustering at the state level is more conservative, as it accounts for arbitrary within-state correlation across mines and over time. The point estimate is unchanged (-0.088) and the state-clustered standard error is 0.021, yielding a t-statistic of 4.2 ($p < 0.001$). With approximately 50 state clusters, this provides reassurance that the results are not driven by a small number of states.

Post-2010 subsample. Restricting the sample to inspections after 2010 yields a point estimate of -0.121 (SE = 0.007). This is somewhat larger than the full-sample estimate, potentially reflecting enhanced information dissemination following the MINER Act of 2006, including improved online access to mine safety records and strengthened miners' rights provisions.

D. Heterogeneity Appendix

Coal versus metal/nonmetal. The coal mining subsample (N = 789,300) produces a point estimate of -0.162 (SE = 0.009), while the metal/nonmetal subsample (N = 3,357,365) produces -0.015 (SE = 0.003). The tenfold difference in magnitude is consistent with (i) higher baseline fatality rates in coal mining, (ii) more extensive S&S findings per inspection in coal mines (which are typically underground and face more numerous hazard categories), and (iii) stronger institutional channels for information transmission in coal mining (higher unionization rates, more formalized safety committees).

Large mines. Restricting to mines with $\geq$20 employees at baseline ($t-1$) yields a sample of 1,018,408 observations. The point estimate of -0.131 (SE = 0.006) exceeds the full-sample estimate, ruling out the concern that results are driven by small mines where single departures generate large percentage swings. Large mines may also exhibit stronger effects because they offer more observable working conditions (larger scale means more opportunities for violations to be noticed by workers) and because workers at large operations have more alternative employment options within the mining sector.

E. Standardized Effect Sizes

Table 5: Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Log employees	-0.0878	0.0043	1.40	-0.0628	0.0031	Moderate negative
Log hours worked	-0.1830	0.0084	2.63	-0.0696	0.0032	Moderate negative

Notes: **Country:** United States. **Research question:** Do workers reduce employment at mines where mandatory safety inspections reveal serious violations, testing the worker information channel of compensating differentials at the establishment level? **Policy mechanism:** MSHA conducts mandatory regular safety inspections of every active mine (4x/year underground, 2x/year surface); inspection results, including Significant & Substantial violation citations, are publicly posted within days, revealing pre-existing hazardous conditions to workers. **Outcome definition:** (Row 1) Log average quarterly employee count from MSHA administrative employment records; (Row 2) Log total quarterly hours worked. **Treatment:** Binary: regular inspection finding $\geq$3 S&S violations ("severe") vs. 0 S&S violations ("clean"). **Data:** MSHA Open Government Data: Inspections (516K regular inspections), Violations (3.06M records), MinesProdQuarterly (2.71M mine-quarter records), 2000–2024. Panel of 4.1M mine-event-quarter observations. **Method:** Stacked event-study DiD with mine and event fixed effects, standard errors clustered at mine level (24,109 mines). **Sample:** All active US mines receiving regular MSHA safety inspections, restricted to events with at least 10 of 13 quarters observable. SDE = $\hat{\beta}$/SD(Y) where SD(Y) is the unconditional standard deviation. Classification refers to magnitude, not statistical significance: Large (|SDE| > 0.15), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).