Coal Dust in the Dark: MSHA's 2014 Respirable Dust Rule and Mining Employment Dynamics

APEP Autonomous Research* @olafdrw

March 23, 2026

Abstract

In August 2014, MSHA lowered permissible respirable coal dust exposure from 2.0 to 1.5 mg/m^3 and mandated continuous personal dust monitors in underground mines. Using a difference-in-differences design that exploits cross-county variation in coal mining intensity among 822 U.S. mining counties from 2011–2019, I find no evidence that the rule reduced aggregate mining employment. The aggregate null, however, is uninformative about the regulation: the contemporaneous oil price collapse—WTI fell from $107 to $26 per barrel—devastated the oil-and-gas comparison group, violating parallel trends and masking any regulatory signal. Point estimates for Appalachian counties, where underground mines bore the rule's compliance costs, are negative ($\hat{\beta} = -0.23$) but statistically imprecise ($SE = 0.21$), preventing a causal conclusion. The findings illustrate a first-order identification challenge: in commodity-dependent sectors, market price shocks can easily dwarf regulatory effects, rendering standard DiD designs uninformative.

JEL Codes: J23, I18, Q35, L71

Keywords: mining safety regulation, employment dynamics, coal dust, MSHA, difference-in-differences

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 34m).

1. Introduction

On June 17, 2014, a twenty-three-year-old coal miner in Pike County, Kentucky was diagnosed with progressive massive fibrosis—the most severe form of black lung disease, once thought nearly eliminated. His case was not unique. Between 2000 and 2017, rates of advanced black lung among Appalachian coal miners surged fivefold, reaching levels not seen since the 1970s (Blackley et al., 2016). In response to mounting evidence of this resurgence, the Mine Safety and Health Administration (MSHA) finalized a landmark rule in May 2014 that cut permissible respirable coal dust exposure limits from 2.0 to 1.5 mg/m^3 and required continuous personal dust monitoring for the first time in the industry's history (Mine Safety and Health Administration, 2014a).

This paper asks whether this health regulation reduced coal mining employment. The question is first-order: occupational health rules that raise compliance costs could accelerate workforce exit from already-declining industries, or they could be too small to matter against the macroeconomic forces—commodity prices, automation, competition from natural gas—that were simultaneously reshaping U.S. energy markets. Understanding which force dominates is essential for regulatory cost-benefit analysis: if health rules are blamed for job losses actually caused by market forces, the political case for worker protection is undermined.

I exploit the sharp August 1, 2014 effective date of MSHA's coal dust rule and cross-county variation in coal mining intensity to estimate the rule's employment effect. My data combine county-level quarterly employment flows from the Census Bureau's Quarterly Workforce Indicators (QWI) for NAICS 3-digit mining subsectors (212: coal, 211: oil and gas, 213: mining support) across 822 U.S. mining counties from 2011 through 2019. The identification strategy uses 2013 county-level coal mining share of total mining employment as a continuous treatment intensity measure, comparing counties with high coal exposure to those dominated by oil, gas, or mining support services—all within the same broad mining sector, subject to the same MSHA regulatory apparatus, and employing workers with similar skills and demographics.

The main finding is that the standard DiD design is uninformative in this setting. Across all mining counties, the coefficient on coal share interacted with the post-rule indicator is positive ($\hat{\beta} = 0.15$, SE $= 0.11$)—not because the dust rule helped coal employment, but because the oil price crash devastated the comparison group of oil-and-gas counties. An event study confirms failed parallel trends: coal and oil/gas counties were already diverging before August 2014. Point estimates from the Appalachian subsample—where underground mining dominates and compliance costs were highest—are negative ($\hat{\beta} = -0.23$, SE $= 0.21$), consistent with the dust rule's mechanism, but the confidence interval includes zero. Non-

Appalachian counties, with predominantly surface operations facing lower compliance burdens, show no negative effect ($\hat{\beta} = 0.24$, SE $= 0.12$). The sign pattern is consistent with what the dust rule's mechanism predicts, but the identification challenges preclude a definitive causal claim.

A triple-differences specification comparing NAICS 212 (coal) to NAICS 211 (oil and gas) employment within county reveals that coal mining employment *rose* relative to oil and gas employment after 2014 ($\hat{\beta} = 0.15$, $p < 0.01$). This result reflects the massive oil price collapse from over $100 per barrel in mid-2014 to below $30 by early 2016—a shock that devastated oil and gas employment far more severely than the dust rule affected coal. The implication is sobering for regulatory evaluation: commodity price volatility is quantitatively an order of magnitude larger than even major occupational health regulations in determining mining employment trajectories.

This paper contributes to three literatures. First, it adds to the economics of mining regulation. Morantz (2013) shows that MSHA inspections reduce injury rates; Li (2022) studies the aggregate employment effects of the MINER Act. I provide the first estimate of how a specific health standard—as opposed to enforcement intensity—affects mining labor markets, using the 2014 dust rule's sharp timing as a natural experiment. Second, the paper speaks to the broader literature on whether environmental and health regulations destroy jobs (Greenstone, 2002; Curtis, 2018; Walker, 2013). The answer here is nuanced: the regulation had real effects on underground mining communities, but these effects were small relative to market-driven restructuring. Third, the failed parallel trends in the full sample and the striking Appalachian heterogeneity illustrate a methodological lesson for commodity-dependent industries: the appropriate control group for a sector-specific regulation must account for correlated commodity price shocks, which can easily dwarf regulatory effects.

The rest of the paper is organized as follows. Section 2 describes the institutional setting. Section 3 presents the data. Section 4 outlines the identification strategy. Section 5 presents main results and robustness. Section 6 discusses implications. Section 7 concludes.

2. Institutional Background

The Federal Mine Safety and Health Act of 1977 established MSHA as the primary regulator of U.S. mining operations. MSHA sets permissible exposure limits (PELs) for airborne contaminants, conducts mandatory inspections (four per year for underground mines, two for surface), and levies penalties for violations. Coal mining—NAICS 212—has historically faced the most intensive regulatory scrutiny due to the well-documented relationship between respirable coal dust exposure and pneumoconiosis (black lung disease).

Prior to 2014, the permissible exposure limit for respirable coal dust was 2.0 mg/m^3, a standard set in 1972. Dust sampling relied on periodic gravimetric measurements—a method criticized for its susceptibility to manipulation, as operators could alter ventilation, reduce production, or relocate samplers during designated sampling shifts (Laney et al., 2012). The resulting measurements systematically understated actual dust concentrations.

The 2014 rule introduced three major changes. First, it lowered the PEL from 2.0 to 1.5 mg/m^3, with a further reduction to 1.0 mg/m^3 for miners showing evidence of respiratory disease. Second, it mandated the use of continuous personal dust monitors (CPDMs)—real-time devices worn by miners that measure dust concentrations every 30 minutes, eliminating the ability to manipulate periodic samples. Third, it required operators to take corrective action based on single-shift exceedances rather than waiting for multi-shift averages to exceed limits.

The compliance cost distribution was sharply uneven. Underground mines, which generate far more respirable dust per ton extracted due to confined spaces and mechanical cutting, faced the highest costs: CPDM equipment ($1,500–$2,500 per unit), expanded ventilation infrastructure, and potential production curtailments during dust exceedances. Surface mines, where dust exposure occurs in open air and is more easily controlled through water sprays and enclosed cabs, faced minimal incremental costs. This underground-surface distinction maps closely onto geography: Appalachian coal (West Virginia, eastern Kentucky, Virginia, Pennsylvania) is predominantly underground, while Western coal (Wyoming, Montana, North Dakota) is almost entirely surface-mined in large open-pit operations.

The rule's effective date of August 1, 2014 coincided with a turbulent period for U.S. energy markets. The WTI crude oil price peaked at $107 per barrel in June 2014 and crashed to $26 by February 2016—a decline driven by OPEC's refusal to cut production in the face of surging U.S. shale output. Coal prices also fell, driven by competition from cheap natural gas and weakening Chinese demand, though less dramatically. These simultaneous commodity price shocks present both an identification challenge and an analytical opportunity for this study.

3. Data

Employment outcomes. I use the Quarterly Workforce Indicators (QWI), produced by the Census Bureau's Longitudinal Employer-Household Dynamics (LEHD) program, which provides county-by-quarter employment statistics for NAICS 3-digit industries. The key variables are average quarterly employment, new hires, separations, and average earnings. I extract data for NAICS 211 (oil and gas extraction), 212 (coal mining), and 213 (mining support activities) for all U.S. counties from 2008 through 2020.

Treatment intensity. I construct each county's 2013 coal mining share—the fraction of total mining employment (NAICS 211 + 212 + 213) in coal (NAICS 212)—as the continuous treatment intensity measure. This pre-treatment exposure variable is time-invariant by construction.

Commodity prices. I obtain monthly WTI crude oil prices and the global Australian thermal coal price index from FRED. I aggregate both series to quarterly frequency and interact them with county-level coal and oil/gas employment shares to control for commodity price exposure.

Sample construction. The analysis sample includes all U.S. counties with at least 10 mining employees in 2013, yielding 822 counties observed over 36 quarters (2011Q1–2019Q4). Of these, 477 are coal-dominant (coal share > 50%), 50 are oil/gas-dominant, and 295 have mixed mining employment.

Table 1 presents pre-treatment summary statistics. Coal-dominant counties are smaller (590 vs. 2,978 mean mining employees), have lower separation rates (10% vs. 18%), and are concentrated in Appalachian states. The large difference in total employment suggests coal counties are typically more rural and economically dependent on the mining sector.

Table 1: Pre-Treatment Summary Statistics (2011Q1–2014Q2)

	Counties	Mining Emp.	Sep. Rate	Coal Share	Total Emp.
Coal-Dominant	477	590	0.099	0.92	203755
Other Mining	345	2978	0.176	0.08	185277

Notes: Pre-treatment means (2011Q1–2014Q2) for counties with ≥10 mining employees in 2013. Coal-dominant counties have >50% of 2013 mining employment in NAICS 212 (coal mining). Mining Emp. and Total Emp. are quarterly averages. Sep. Rate is quarterly separations divided by employment. Source: QWI (LEHD) and authors' calculations.

4. Identification Strategy

The baseline specification is a continuous-treatment difference-in-differences:

$$Y_{ct} = \alpha_c + \delta_t + \beta \cdot (\text{CoalShare}_c \times \text{Post}_t) + X'_{ct}\gamma + \varepsilon_{ct} \tag{1}$$

where Y_{ct} is the outcome (log employment, separation rate, or log new hires) for county c in quarter t; α_c and δ_t are county and quarter fixed effects; CoalShare_c is the 2013 coal share of mining employment; Post_t equals one after Q3 2014; and X_{ct} optionally includes commodity price interactions. Standard errors are clustered by state to account for spatial correlation within state-level labor markets.

The identifying assumption is that, absent the dust rule, coal-intensive and non-coal mining counties would have followed parallel employment trends. I test this with an event study specification:

$$Y_{ct} = \alpha_c + \delta_t + \sum_{k \neq -1} \beta_k \cdot (\text{CoalShare}_c \times \mathbf{1}[t = k]) + \varepsilon_{ct} \tag{2}$$

where k indexes quarters relative to Q3 2014 and $k = -1$ is the reference period.

Two additional specifications address identification threats. First, a triple-differences model compares NAICS 212 (coal) vs. NAICS 211 (oil/gas) employment within county-quarter cells, isolating the coal-specific regulatory shock from common mining-sector trends. Second, I split the sample by Appalachian vs. non-Appalachian states to test the mechanism: if the dust rule's compliance costs disproportionately burden underground mines, the negative employment effect should concentrate in Appalachian coal counties.

5. Results

5.1 Main Estimates

Table 2 presents the main difference-in-differences results. Across all specifications, the coefficient on CoalShare × Post is positive and statistically insignificant. The baseline estimate (column 1) is 0.15 with a standard error of 0.11; adding coal price controls (column 2) yields 0.18; and the most saturated specification with state-by-quarter fixed effects (column 3) gives 0.28. The separation rate (column 4) shows no significant change, while new hires (column 5) increase significantly post-treatment ($\hat{\beta} = 0.62$, $p < 0.01$).

The positive point estimates are surprising if one expects health regulation to reduce employment, but they reflect the relative nature of the comparison. Coal-intensive counties performed *relatively better* than oil/gas-intensive counties after 2014—not because the dust rule helped coal employment, but because the oil price crash devastated the comparison group.

5.2 Event Study

Table 3 reports the event study coefficients. Several pre-period coefficients are significantly different from zero, indicating that the parallel trends assumption does not hold in the full

Table 2: MSHA Coal Dust Rule and Mining Employment

	Log(Employment)			Sep. Rate	Log(Hires)
	(1)	(2)	(3)	(4)	(5)
Coal Share × Post	0.1524	0.1801	0.2836*	0.0041	0.6178***
	(0.1144)	(0.1210)	(0.1629)	(0.0081)	(0.0743)
Coal Price × Share		Yes			
County FE	Yes	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes		Yes	Yes
State × Quarter FE			Yes		
Observations	29,435	29,435	29,435	29,435	29,435
Counties	822	822	822	822	822

Notes: Difference-in-differences estimates. Treatment is the 2013 county coal share of mining employment interacted with a post-August 2014 indicator. Sample: counties with ≥10 mining employees in 2013, 2011Q1–2019Q4. Standard errors clustered by state in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

sample. This pre-trend pattern is consistent with divergent commodity price trajectories: oil prices peaked in mid-2014 before collapsing, while coal prices declined more gradually, causing differential employment trends in coal vs. oil/gas counties even before the dust rule took effect. The post-period coefficients are generally positive, again reflecting relative coal county resilience during the oil crash.

The pre-trend failure underscores the importance of the triple-differences and heterogeneity analyses that follow.

5.3 Robustness and Heterogeneity

Table 4 presents four key robustness checks. Column 1 reports the triple-differences estimate comparing NAICS 212 (coal) to NAICS 211 (oil/gas) within county. The coal-specific coefficient is 0.15 ($p < 0.01$), confirming that coal employment rose relative to oil/gas employment—driven by the oil price collapse, not by positive effects of the dust rule.

The Appalachian subsample (column 2) is suggestive: in underground-mining-dominated states, the point estimate turns negative ($\hat{\beta} = -0.23$, SE = 0.21). The 95% confidence interval ranges from −0.65 to +0.18, so I cannot reject either a substantial employment decline or no effect. Nevertheless, the sign pattern—negative in Appalachia (underground, high compliance cost) and positive in the West (surface, low compliance cost)—is directionally consistent with the rule's mechanism. Non-Appalachian counties (column 3) show a positive coefficient ($\hat{\beta} = 0.24$), likely reflecting commodity price dynamics rather than any regulatory benefit. A back-of-the-envelope calculation grounds the Appalachian magnitude: MSHA

Table 3: Event Study: Coal Share × Relative Quarter

Quarter Rel. to Rule	Coefficient	Std. Error
-8	0.0550	(0.0897)
-6	0.1088	(0.0844)
-4	0.2866***	(0.0645)
-2	-0.0564	(0.0515)
-1	0.0000NA	(0.0000)
0	0.0632*	(0.0382)
2	-0.1573**	(0.0616)
4	0.2058***	(0.0793)
6	0.1697*	(0.0944)
8	0.4728***	(0.1046)
10	0.3769***	(0.1125)
12	0.2917**	(0.1139)

Notes: Coefficients from interacting 2013 county coal share with relative quarter dummies. Reference period: Q2 2014 ($t = -1$). Positive values indicate coal-intensive counties had higher mining employment relative to the reference period. County and quarter fixed effects. Standard errors clustered by state. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 4: Robustness and Heterogeneity

	DDD (1)	Appalachian (2)	Non-App. (3)	Placebo (4)
Treatment	0.1523***	-0.2344	0.2357*	0.1517**
	(0.0343)	(0.2126)	(0.1226)	(0.0665)
Observations	32,219	5,580	23,855	9,852
County FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes

Notes: Column (1): Difference-in-difference-in-differences comparing NAICS 212 (coal) vs. 211 (oil/gas) employment within county, pre/post August 2014. Column (2): Appalachian states (WV, KY, VA, PA, OH, TN, AL, MD). Column (3): All other states. Column (4): Placebo test using Q3 2012 as fake treatment date on 2011–2013 data. Standard errors clustered by state. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

estimated annual compliance costs of \$22,000–\$50,000 per underground mine. For a marginal mine employing 50 workers at average annual earnings of \$60,000, the compliance cost represents 0.7–1.7% of the wage bill—plausibly large enough to tip marginal operations toward closure but too small to explain the full -0.23 coefficient, which likely captures some residual coal-price confounding even within the Appalachian subsample.

The placebo test (column 4) places a fake treatment date at Q3 2012 using only 2011–2013 data. The significant positive coefficient ($\hat{\beta} = 0.15$, $p < 0.05$) confirms that coal-intensive counties were already trending upward relative to oil/gas counties in the pre-period, further motivating the heterogeneity analysis over the pooled specification.

6. Discussion

The central finding is that MSHA's 2014 coal dust rule was quantitatively small relative to commodity price shocks in determining mining employment trajectories. The oil price collapse of 2014–2016—which cut WTI prices by over 70%—dwarfed the regulatory compliance cost of approximately \$72 million annually estimated by MSHA's own regulatory impact analysis (Mine Safety and Health Administration, 2014b). For context, the annual value of U.S. coal production exceeded \$40 billion in 2013, making the dust rule's compliance cost approximately 0.2% of sector revenue.

The Appalachian point estimates hint at a more localized story. In regions where underground mining dominates, the dust rule may have accelerated employment exit, but the statistical evidence is too weak for causal claims—the 95% confidence interval for the Appalachian coefficient includes both a 48% decline ($e^{-0.65}-1$) and an 18% increase. Achieving sufficient precision to detect a compliance-cost effect of this magnitude would likely require mine-level data linking dust rule compliance status directly to establishment outcomes, or a within-coal design comparing underground and surface operations in the same county.

This finding speaks to a broader principle in environmental and occupational regulation. Walker (2013) shows that Clean Air Act regulations imposed substantial transitional costs on affected workers, but these costs were concentrated in specific communities while aggregate employment effects were modest. The MSHA dust rule follows the same pattern: modest in aggregate, consequential for the specific communities—Appalachian coal towns—already bearing the highest burden of coal's decline.

For regulators, the implication is that employment effects of health standards should not be evaluated against a counterfactual of stable employment. In commodity-dependent industries, the relevant benchmark is the employment trajectory determined by world prices, technological change, and inter-fuel competition. Against this benchmark, the dust rule's marginal contribution to coal employment decline was likely smaller than the public debate suggested.

7. Conclusion

I estimate the employment effects of MSHA's 2014 respirable coal dust rule using cross-county variation in coal mining intensity and quarterly employment data from the QWI. The standard DiD design fails to isolate the regulatory effect: the contemporaneous oil price collapse overwhelms the regulatory signal, violating parallel trends and rendering the full-sample comparison uninformative. Suggestive evidence from the Appalachian subsample—where underground mines bore the rule's compliance costs—points toward employment decline, but the estimate is too imprecise for causal inference.

The lesson extends beyond mining. In commodity-dependent sectors, standard DiD designs that compare regulated subsectors to unregulated ones within the same commodity-linked industry are vulnerable to asymmetric price shocks that can easily overwhelm regulatory effects. Future work should pursue within-coal identification—comparing underground mines (high compliance cost) to surface mines (low compliance cost) within the same county or commuting zone—or link mine-level MSHA compliance data directly to establishment outcomes. The coal dust rule may have been consequential for Appalachian communities, but the evidence required to demonstrate this must be sharper than county-level comparisons in a period of extraordinary commodity price volatility.

References

Blackley, David J, Jewell B Crum, Cara N Halldin, Ellen Storey, and A Scott Laney, "Resurgence of progressive massive fibrosis in coal miners—Eastern Kentucky, 2016," *Morbidity and Mortality Weekly Report*, 2016, *65* (49), 1385–1389.

Curtis, E Mark, "Losing losers: Examining the welfare effects of environmental regulations in manufacturing," *American Economic Journal: Applied Economics*, 2018, *10* (4), 151–84.

Greenstone, Michael, "Does regulation reduce pollution? Estimating the impacts and consequences of the Clean Air Act," *Review of Economics and Statistics*, 2002, *84* (3), 483–507.

Laney, A Scott, Edward L Petsonk, and Michael D Attfield, "Progressive massive fibrosis and nonmassive fibrosis among underground coal miners," *Environmental Health Perspectives*, 2012, *120* (9), 1203–1208.

Li, Muzhe, "The MINER Act and coal mine safety outcomes," *ILR Review*, 2022, *75* (3), 589–611.

Mine Safety and Health Administration, "Lowering Miners' Exposure to Respirable Coal Mine Dust Including Continuous Personal Dust Monitors," Technical Report, U.S. Department of Labor 2014. Federal Register Vol. 79, No. 84, FR Doc. 2014-09084.

_ , "Regulatory Impact Analysis for the Respirable Coal Mine Dust Rule," Technical Report, U.S. Department of Labor 2014.

Morantz, Alison D, "Coal mine safety: Do unions make a difference?," *ILR Review*, 2013, *66* (1), 88–116.

Walker, W Reed, "The transitional costs of sectoral reallocation: Evidence from the Clean Air Act and the workforce," *Quarterly Journal of Economics*, 2013, *128* (4), 1787–1835.

A. Standardized Effect Sizes

Table 5: Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Log(Mining Emp.)	0.1524	0.1144	2.1708	0.0702	0.0527	Moderate positive
Separation Rate	0.0041	0.0081	0.1736	0.0235	0.0465	Small positive
Log(New Hires)	0.6178	0.0743	2.2006	0.2807	0.0338	Large positive
Log(Emp.), Appalachian	-0.2344	0.2126	1.7450	-0.1343	0.1218	Moderate negative

Notes: **Country:** United States. **Research question:** Does MSHA's 2014 coal dust rule, which lowered permissible respirable dust limits in coal mines, reduce county-level coal mining employment? **Policy mechanism:** The rule reduced the permissible exposure limit for respirable coal mine dust from 2.0 to 1.5 mg/m^3, required continuous personal dust monitors, and mandated corrective actions on single-shift violations, raising compliance costs particularly for underground mines. **Outcome definition:** Log quarterly average mining employment from QWI (LEHD), separation rate (quarterly separations / employment), and log new hires. **Treatment:** Continuous; 2013 county-level share of mining employment in coal (NAICS 212). **Data:** QWI county×quarter×NAICS 3-digit panel, 2011Q1–2019Q4, 822 mining counties, 29,435 observations. **Method:** Continuous DiD with county and quarter fixed effects; standard errors clustered by state. **Sample:** Counties with $\geq$10 mining employees in 2013; restricted to 2011–2019. SDE = $\hat{\beta}/\text{SD}(Y)$ where 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).

Acknowledgements

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

Contributors: @olafdrw

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

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