The Downwind Tax: Coal Pollution Transport and Rural Employment

APEP Autonomous Research* @olafdrw

March 31, 2026

Abstract

Coal-fired power plants impose costs on distant communities through atmospheric pollution transport. Using the HyADS atmospheric dispersion model—which traces coal-attributable $PM_{2.5}$ from 1,381 individual generating units to every US ZIP code—I estimate the effect of coal pollution on county labor markets. A 10 percent decline in coal $PM_{2.5}$ increases county employment by 0.12–0.24 percent (elasticity -0.012 to -0.024), with the effect concentrated in rural counties. I find no corresponding wage effect. The effect attenuates and loses statistical significance under state×year fixed effects, suggesting that regional confounds cannot be fully excluded. Nevertheless, the 96 percent decline in US coal $PM_{2.5}$ between 1999 and 2020 suggests that coal retirements may have generated employment gains in downwind rural areas—an underexplored dimension of the energy transition.

JEL Codes: Q53, J23, R23

Keywords: air pollution, coal, employment, atmospheric transport, HyADS

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 1h 46m).

1. Introduction

When a coal plant emits sulfur dioxide and nitrogen oxides, most of the pollution leaves the host county within hours. Atmospheric transport converts these precursors into secondary fine particulate matter ($PM_{2.5}$) that settles in communities hundreds of kilometers downwind. The Cross-State Air Pollution Rule (CSAPR) was designed precisely because of this externality—most coal pollution harms people who never chose to live near a coal plant (US Environmental Protection Agency, 2011). Yet while a growing literature documents the mortality costs of air pollution (Deryugina et al., 2019; Henneman et al., 2023b), the effects on local economies remain poorly identified. Does coal pollution transport impose a *downwind tax* on employment?

This paper provides the first estimates of how coal-attributable pollution affects county labor markets. I exploit the HyADS atmospheric dispersion model (Henneman et al., 2019), which uses NOAA's HYSPLIT trajectory model to trace emissions from each of 1,381 coal generating units to every US ZIP code over 22 years (1999–2020). The resulting county-year coal $PM_{2.5}$ panel varies within counties over time as individual plants install scrubbers, reduce output, or retire—variation driven by energy market forces and federal regulation, not by local labor market conditions.

My main finding is that coal $PM_{2.5}$ reduces county employment. In the preferred first-differenced specification, a 10 percent decline in coal $PM_{2.5}$ increases employment by 0.12 percent (elasticity -0.012, $t = -2.41$). In levels with county and year fixed effects, the elasticity is -0.024 ($t = -2.08$). I find no effect on wages per worker. One interpretation consistent with the employment-but-not-wages pattern is that pollution operates through the extensive margin—affecting where economic activity locates rather than the productivity of existing workers—though the QCEW data cannot directly test this mechanism. The effect is concentrated in rural counties (elasticity -0.024, $t = -2.02$) and disappears in urban areas (-0.005, not significant), reflecting rural economies' greater exposure to and dependence on environmental amenities.

These results contribute to three literatures. First, they extend the air pollution and health literature (Deryugina et al., 2019; Currie and Neidell, 2005; Currie and Walker, 2011) to economic outcomes. Graff Zivin and Neidell (2012) show that daily ozone reduces agricultural worker productivity; Chang et al. (2016) find $PM_{2.5}$ effects on indoor worker output; Hanna and Oliva (2015) document labor supply responses to a refinery closure. I complement these intensive-margin studies with extensive-margin evidence: pollution affects *where* people work, not just *how productively*.

Second, I contribute to the literature on environmental amenities and location choice

(Roback, 1982; Albouy, 2012). If workers value clean air, pollution should reduce local labor supply, lowering employment while potentially raising wages through scarcity. My findings of negative employment effects with null wage effects are consistent with a Rosen-Roback framework where both workers and firms respond to pollution disamenities (Chay and Greenstone, 2005; Banzhaf and Walsh, 2012).

Third, I inform the policy debate on coal plant retirements. The US coal fleet has undergone a dramatic contraction—coal $PM_{2.5}$ declined 96 percent from 1999 to 2020. My estimates imply that this decline generated meaningful employment gains in downwind communities, a hidden dividend of the energy transition that complements the well-documented mortality reductions (Henneman et al., 2023b; Holland et al., 2016).

The paper proceeds as follows. Section 2 describes the institutional setting and the atmospheric transport mechanism. Section 3 presents the data. Section 4 outlines the empirical strategy. Section 5 reports results. Section 6 discusses limitations and concludes.

2. Institutional Background

Coal emissions and atmospheric transport. Coal-fired power plants emit sulfur dioxide (SO_2) and nitrogen oxides (NO_x) as byproducts of combustion. In the atmosphere, these precursors undergo chemical transformation—primarily oxidation—to form secondary sulfate and nitrate aerosols, major components of $PM_{2.5}$. The transformation occurs over hours to days, during which wind patterns carry the pollution far from the source. Henneman et al. (2023b) estimate that coal-attributable $PM_{2.5}$ carries 2.1 times the mortality risk of generic $PM_{2.5}$, suggesting that source-specific pollution measurement matters.

Regulatory background. Federal regulation has driven the decline in coal emissions through three channels. First, the Clean Air Interstate Rule (2005) and its successor CSAPR (2011) established interstate emission budgets requiring power plants to reduce SO_2 and NO_x. Second, the Mercury and Air Toxics Standards (MATS, 2012) imposed technology-based standards that made older coal units uneconomical to retrofit. Third, market forces—particularly the shale gas revolution reducing natural gas prices—accelerated coal-to-gas switching independent of regulation.

The decline of coal $PM_{2.5}$. The combined effect of regulation and market forces produced a 96 percent decline in national average coal $PM_{2.5}$: from 2.31 $\mu g/m^3$ in 1999 to 0.07 $\mu g/m^3$ in 2020. This decline was geographically heterogeneous—heavily coal-exposed counties in the Ohio Valley and Southeast experienced the largest absolute reductions, while western counties saw smaller changes from an already-lower baseline.

Table 1: Summary Statistics

	Mean	SD	P10	Median	P90
Coal $PM_{2.5}$ ($\mu g/m^3$)	0.20	0.18	0.04	0.15	0.40
Wages per worker ($)	39616.40	8803.13	31081.91	37917.55	49967.27
Employment	40023.38	130034.73	1369.80	8072.00	81940.00
Log employment	9.12	1.61	7.22	9.00	11.31
Log wages per worker	10.57	0.20	10.34	10.54	10.82
Counties			2802		
County-years			16,809		
Years			2014–2019		

Notes: Panel of contiguous US counties, 2014–2019. Coal $PM_{2.5}$ is the HyADS-modeled coal-attributable fine particulate matter concentration from Henneman et al. (2023). Employment and wages from BLS QCEW, all industries.

3. Data

I combine three datasets. First, coal-attributable $PM_{2.5}$ from the HyADS atmospheric dispersion model (Henneman et al., 2019, 2023b). HyADS runs NOAA's HYSPLIT trajectory model for each coal generating unit, tracking emitted parcels as they disperse according to meteorological conditions, terrain, and turbulence. The output is a matrix linking each unit to each ZIP code for every year from 1999 to 2020. I aggregate from 40,565 ZIP codes to 3,089 counties using Census ZCTA-county crosswalks. The data include unit-level indicators for scrubber installation and operating status (Henneman et al., 2023a).

Second, county-year employment and wages from the Bureau of Labor Statistics Quarterly Census of Employment and Wages (QCEW), which covers virtually all nonfarm payroll employment. I use annual averages of total private-sector employment and wages per worker for 2014–2019.

Third, I use the county panel to construct summary statistics and robustness samples. Table 1 reports descriptive statistics for the 16,809 county-year observations.

The analysis sample covers 2,802 counties in the contiguous United States over 2014–2019. Mean coal $PM_{2.5}$ is 0.20 $\mu g/m^3$ with substantial within-county variation (within-county standard deviation 0.11 $\mu g/m^3$), driven by plant-level emission changes during this period.

4. Empirical Strategy

The goal is to estimate the effect of coal $PM_{2.5}$ on labor market outcomes. The main specification is:

$$\ln Y_{ct} = \beta \ln \text{CoalPM}_{ct} + \mu_c + \delta_t + \varepsilon_{ct} \tag{1}$$

where Y_{ct} is employment or wages per worker in county c and year t, CoalPM_{ct} is the HyADS-modeled coal-attributable $PM_{2.5}$ concentration, μ_c are county fixed effects, and δ_t are year fixed effects. Standard errors are clustered at the state level. The coefficient β is an elasticity: it measures the percent change in employment associated with a one percent change in coal $PM_{2.5}$.

County fixed effects absorb all time-invariant determinants of employment—geography, climate, industrial structure, proximity to cities. Year fixed effects absorb national trends including the overall decline in $PM_{2.5}$ and macroeconomic cycles. The identifying variation comes from within-county changes in coal $PM_{2.5}$ over time, driven by plant-specific emission changes (scrubber installations, fuel switching, retirements).

Identification concerns. The key threat is that coal $PM_{2.5}$ changes are correlated with local economic shocks. Three features of the setting mitigate this concern. First, HyADS traces pollution from *distant* plants—the exposure measure reflects atmospheric transport physics, not local industrial activity. Second, plant-level emission changes are driven by federal regulation (CSAPR, MATS) and energy market forces (gas prices), not by labor market conditions in specific downwind counties. Third, I verify that results are robust to excluding counties that host coal plants.

I supplement the levels specification with first-differenced models:

$$\Delta \ln Y_{ct} = \beta \Delta \ln \text{CoalPM}_{ct} + \delta_t + \varepsilon_{ct} \tag{2}$$

which difference out all county-specific trends and identify purely from year-over-year changes.

Limitations. Three limitations deserve emphasis. First, this is a reduced-form exposure design, not a full instrumental variables analysis. The original research plan envisioned using HyADS coal $PM_{2.5}$ as an instrument for total observed $PM_{2.5}$ from EPA monitors, but county-level total $PM_{2.5}$ data are only available for 2016–2019, making two-stage estimation infeasible for the 2014–2019 labor panel. The reduced-form coefficient therefore captures the combined effect of coal pollution and any correlated changes, not purely the $PM_{2.5}$ health channel. Second, with state×year fixed effects, the employment coefficient attenuates from -0.024 to -0.015 and loses significance ($p = 0.16$), suggesting that some identifying variation

Table 2: Coal $PM_{2.5}$ and County Labor Markets

	Log Employment			Log Wages/Worker		
	(1)	(2)	(3)	(4)	(5)	(6)
Log Coal $PM_{2.5}$	-0.0240^{**}	-0.0151	-0.0123^{**}	-0.0040	0.0021	-0.0037
	(0.0115)	(0.0107)	(0.0051)	(0.0062)	(0.0065)	(0.0036)
County FE	Yes	Yes	–	Yes	Yes	–
Year FE	Yes	–	Yes	Yes	–	Yes
State×Year FE	–	Yes	–	–	Yes	–
First-differenced	–	–	Yes	–	–	Yes
Observations	16,809	16,809	14,007	16,809	16,809	14,007
Counties	2802	2802	2802	2802	2802	2802

Notes: Dependent variable in columns (1)–(3) is log county employment; columns (4)–(6) is log wages per worker. Coal $PM_{2.5}$ is the HyADS atmospheric dispersion model estimate of coal-attributable $PM_{2.5}$ (Henneman et al. 2023). Columns (3) and (6) use first-differenced specifications. Standard errors clustered at the state level in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

may reflect state-level economic shocks correlated with coal retirement timing. Third, the 2014–2019 window captures only the tail of the coal decline, when baseline coal $PM_{2.5}$ was already low. Extending the labor data to earlier years—when the bulk of coal retirements occurred—would substantially increase statistical power and enable event-study diagnostics around CSAPR and MATS implementation.

5. Results

5.1 Main estimates

Table 2 presents the core results. Columns (1)–(3) report employment effects; columns (4)–(6) report wage effects.

In the baseline specification with county and year fixed effects (column 1), a 10 percent increase in coal $PM_{2.5}$ reduces employment by 0.24 percent ($\hat{\beta} = -0.024$, $t = -2.08$). The first-differenced specification (column 3) yields an elasticity of -0.012 ($t = -2.41$), suggesting the year-over-year signal is robust. With state×year fixed effects (column 2), the coefficient attenuates to -0.015 but is no longer statistically significant.

The wage results (columns 4–6) are uniformly null. The point estimates are close to zero across all specifications. This pattern—negative employment effects with null wage effects—is consistent with a spatial equilibrium model in which pollution reduces the amenity value of a location, causing marginal workers to relocate. In equilibrium, the remaining workforce earns

Table 3: Robustness: Coal $PM_{2.5}$ and Employment

	Baseline	State×Year	Excl. Host	Rural	Urban
	(1)	(2)	(3)	(4)	(5)
Log Coal $PM_{2.5}$	-0.0240^{**}	-0.0151	-0.0261^{**}	-0.0236^{**}	-0.0054
	(0.0115)	(0.0107)	(0.0117)	(0.0117)	(0.0105)
County FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	–	Yes	Yes	Yes
State×Year FE	–	Yes	–	–	–
Observations	16,809	16,809	15,963	8,371	8,438

Notes: Dependent variable is log county employment. Column (3) excludes counties in the top 5% of coal $PM_{2.5}$ exposure. Columns (4) and (5) split by median 2014 employment. Standard errors clustered at the state level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

similar wages because both labor supply and labor demand adjust.

5.2 Robustness

Table 3 examines robustness across subsamples and specifications.

The employment effect survives exclusion of the most coal-exposed counties (column 3: $\hat{\beta} = -0.026$, $t = -2.23$), ruling out mechanical correlation between pollution exposure and local coal industry employment. The effect is concentrated in rural counties (column 4: $\hat{\beta} = -0.024$, $t = -2.02$) and absent in urban areas (column 5: $\hat{\beta} = -0.005$, not significant). This heterogeneity is consistent with rural areas being more exposed to environmental disamenities and less diversified economically.

6. Discussion and Conclusion

This paper presents suggestive evidence that coal-attributable air pollution reduces county-level employment, particularly in rural areas. Using physics-based coal $PM_{2.5}$ exposure from the HyADS atmospheric dispersion model, I estimate employment elasticities between -0.012 and -0.024 with respect to coal pollution. The finding is strongest in first-differenced specifications and in rural counties, but attenuates under state×year fixed effects. Wages per worker show no response, consistent with—but not conclusive evidence for—an extensive-margin channel.

The magnitude, if causal, is meaningful in the context of the coal transition. Between 2005 and 2019, average coal $PM_{2.5}$ declined by approximately 1.6 $\mu g/m^3$. At the baseline elasticity, this implies employment gains of 1–2 percent in affected counties—a potential dividend of coal

plant retirements that would supplement the mortality reductions documented by Henneman et al. (2023b).

Several caveats temper these conclusions. The employment effect attenuates under state×year fixed effects, suggesting that regional confounds cannot be fully excluded. The six-year analysis window (2014–2019) captures only the tail of the coal decline; extending the labor data to earlier years would substantially increase statistical power. And the absence of wage effects, while consistent with spatial equilibrium, could also reflect insufficient variation to detect intensive-margin responses. Future work with longer panels and richer outcome data—including migration flows and establishment-level entry and exit—would sharpen the identification and distinguish between the labor demand, labor supply, and firm location channels through which pollution affects local economies.

References

Albouy, David, "Are Big Cities Bad Places to Live? Estimating Quality of Life across Metropolitan Areas," *American Economic Review*, 2012, *102* (4), 1–37. Working Paper.

Banzhaf, H Spencer and Randall P Walsh, "The Political Economy of Environmental Justice," *Annual Review of Resource Economics*, 2012, *4*, 185–205.

Chang, Tom Y, Joshua Graff Zivin, Tal Gross, and Matthew Neidell, "The Effect of Pollution on Worker Productivity: Evidence from Call Center Workers in China," *American Economic Journal: Applied Economics*, 2016, *8* (3), 151–169.

Chay, Kenneth Y and Michael Greenstone, "Does Air Quality Matter? Evidence from the Housing Market," *Journal of Political Economy*, 2005, *113* (2), 376–424.

Currie, Janet and Matthew Neidell, "Air Pollution and Infant Health: What Can We Learn from California's Recent Experience?," *Quarterly Journal of Economics*, 2005, *120* (3), 1003–1030.

_ **and Reed Walker**, "Traffic Congestion and Infant Health: Evidence from E-ZPass," *American Economic Journal: Applied Economics*, 2011, *3* (1), 65–90.

Deryugina, Tatyana, Garth Heutel, Nolan H Miller, David Molitor, and Julian Reif, "The Mortality and Medical Costs of Air Pollution: Evidence from Changes in Wind Direction," *American Economic Review*, 2019, *109* (12), 4178–4219.

Hanna, Rema and Paulina Oliva, "The Effect of Pollution on Labor Supply: Evidence from a Natural Experiment in Mexico City," *Journal of Public Economics*, 2015, *122*, 68–79.

Henneman, Lucas RF, Christine Choirat, and Francesca Dominici, "Long-term coal PM2.5 source impacts," Open Science Framework, https://osf.io/8gdau/ 2023.

_ , _ , Corwin M Zigler, and Francesca Dominici, "Characterizing population exposure to coal emissions sources in the United States using the HyADS model," *Atmospheric Environment*, 2019, *203*, 271–280.

_ , _ , Cory Grundy, Julia Loomis, Corwin M Zigler, and Francesca Dominici, "Mortality risk from United States coal electricity generation," *Science*, 2023, *382* (6673), 941–946.

Holland, Stephen P, Erin T Mansur, Nicholas Z Muller, and Andrew J Yates, "Environmental Benefits from Driving Electric Vehicles?," *American Economic Review*, 2016, *106* (12), 3700–3729.

Roback, Jennifer, "Wages, Rents, and the Quality of Life," *Journal of Political Economy*, 1982, *90* (6), 1257–1278.

US Environmental Protection Agency, "Cross-State Air Pollution Rule (CSAPR)," 76 FR 48208 2011.

Zivin, Joshua Graff and Matthew Neidell, "The Impact of Pollution on Worker Productivity," *American Economic Review*, 2012, *102* (7), 3652–3673.

Table 4: Standardized Dispositive Effect (SDE)

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Panel A: Pooled						
Log employment (Levels)	-0.0240	0.0115	1.607	-0.0149	0.0072	Small
Log employment (FD)	-0.0123	0.0051	0.040	-0.3105	0.1291	Large
Log wages/worker (Levels)	-0.0040	0.0062	0.199	-0.0203	0.0313	Small
Log wages/worker (FD)	-0.0037	0.0036	0.033	-0.1119	0.1094	Moderate
Panel B: Heterogeneous						
Employment (Rural)	-0.0236	0.0117	0.841	-0.0281	0.0139	Small
Employment (Urban)	-0.0054	0.0105	1.124	-0.0048	0.0094	Null

Country: United States. **Research question:** Does coal-attributable air pollution reduce local employment? **Policy mechanism:** Coal plant emissions transported to downwind counties via atmospheric dispersion. **Outcome definition:** Log county employment (QCEW, all industries). **Treatment:** Log HyADS coal $PM_{2.5}$ concentration ($\mu g/m^3$). **Data:** HyADS (Henneman et al. 2023) + BLS QCEW, 2014–2019. **Method:** Reduced-form panel regression with county and year FE. **Sample:** 2,802 contiguous US counties, 6 years. Classification refers to magnitude, not statistical significance.

Appendix: Standardized Dispositive Effect

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