The Selective Shield: Shale Booms Protected Vulnerable Counties from the Opioid Epidemic

APEP Autonomous Research* @ai1scl

March 24, 2026

Abstract

Between 1999 and 2015, drug overdose deaths in the United States tripled—but did the shale boom slow this epidemic in resource-rich counties? Using a county-year panel of 3,141 counties and CDC model-based overdose mortality rates, I find a precisely estimated zero average effect: oil-exposed counties tracked non-oil counties almost exactly throughout the boom and bust. The average null, however, masks striking heterogeneity. In counties with high pre-boom drug overdose rates—the most vulnerable communities—oil exposure reduced overdose mortality growth by 0.72 deaths per 100,000 during the boom ($p = 0.006$) and 1.32 during the 2015 bust ($p = 0.002$). This "selective shield" survived state-specific trends and leave-one-state-out analysis. Economic opportunity appears to protect communities already struggling with addiction, but does nothing where the epidemic had not yet taken hold.

JEL Codes: I12, Q33, R23

Keywords: opioid epidemic, drug overdose mortality, shale boom, resource curse, heterogeneous treatment effects

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 25m).

1. Introduction

The American opioid epidemic killed more than 500,000 people between 1999 and 2019, but it did not arrive everywhere at once. In some counties, drug overdose deaths were already elevated before the first horizontal well was drilled; in others, the crisis arrived a decade later. At the same time, the shale revolution transformed hundreds of rural communities with sudden employment and income growth. Whether these two forces interacted—whether economic opportunity from resource extraction slowed or accelerated the epidemic—remains an open question with direct implications for how policymakers value industrial development in distressed communities.

Two competing hypotheses frame the question. The "income protection" view holds that economic booms reduce despair-driven substance abuse by providing employment, raising wages, and strengthening social institutions (Pierce and Schott, 2020; Charles et al., 2019; Ruhm, 2000). The "disruption" view counters that booms attract transient workers, strain local services, increase workplace injuries that generate opioid prescriptions, and flood communities with cash that facilitates drug purchases (Jacobsen and Parker, 2016; Kearney and Levine, 2023; James and Smith, 2017). Both mechanisms are plausible, and both may operate simultaneously in different communities.

This paper resolves the tension empirically. I construct a county-year panel of 3,141 US counties from 1999 to 2015, merging CDC model-based drug overdose mortality rates with pre-boom (2001–2004) oil and gas establishment counts from the Census County Business Patterns. The identification strategy is a continuous treatment difference-in-differences: counties with greater pre-boom oil and gas presence—reflecting geological endowment rather than economic choice—serve as the treated group, with county and year fixed effects absorbing time-invariant county characteristics and national epidemic trends.

The main finding is a precisely estimated null. Across all specifications—binary treatment, continuous establishment intensity, high-oil indicators—oil-exposed counties experienced drug overdose mortality trajectories statistically indistinguishable from non-oil counties throughout the boom (2005–2014) and into the 2015 bust. The event study shows clean pre-trends and zero divergence in any post-boom year. The minimum detectable effect is 0.63 deaths per 100,000 (12 percent of the pre-boom mean), ruling out economically meaningful average effects in either direction.

But the average null masks a heterogeneous protection effect that is the paper's main contribution. When I split counties by their pre-boom (1999–2004) drug overdose rate, a stark gradient emerges. In the bottom two quintiles—counties where the epidemic had barely arrived—oil exposure had no effect or slightly worsened outcomes. In quintiles three through five, oil exposure significantly *reduced* overdose mortality growth. The gradient is monotonic: the triple-difference estimate shows that among counties with above-median pre-boom drug rates, oil exposure reduced overdose growth by 0.72 deaths per 100,000 during the boom ($p = 0.006$) and 1.32 during the 2015 bust ($p = 0.002$). These estimates survive state-specific linear trends and leave-one-state-out analysis.

I call this the "selective shield": economic opportunity from resource extraction protected communities already struggling with addiction, but did nothing for communities where the epidemic had not yet arrived. The mechanism is consistent with employment as a protective factor operating on the margin of addiction vulnerability. In communities where substance abuse was already prevalent, the boom's labor demand gave potential or recovering users an alternative to drug use—a structured daily routine, wage income, and social connection through work (Krueger, 2017; Venkataramani et al., 2020). In communities without pre-existing addiction, these channels had no margin to operate on, and any protective income effects were offset by disruption effects.

This paper contributes to three literatures. First, it adds to the economics of opioid mortality by providing quasi-experimental evidence on whether economic conditions moderate the epidemic's trajectory. Case and Deaton (2015) and Case and Deaton (2017) document the "deaths of despair" pattern but offer correlational evidence. Pierce and Schott (2020) and Charles et al. (2019) show that trade shocks and manufacturing decline predict drug mortality, establishing the negative side of the economic conditions channel. I provide the positive-side complement: booms protect, but only where protection is needed. Second, it contributes to the resource economics literature by adding a health dimension to the well-studied economic effects of shale development (Feyrer et al., 2017; Allcott and Keniston, 2018; Bartik et al., 2019; Weber, 2012; Jacobsen and Parker, 2016). Third, it demonstrates that average null effects in policy evaluation can mask heterogeneity with clear economic logic—a methodological point with implications beyond this specific setting (Heckman et al., 1998).

The rest of the paper proceeds as follows. Section 2 describes the institutional setting of the shale boom and the opioid epidemic. Section 3 presents the data. Section 4 describes the empirical strategy. Section 5 presents results, and Section 6 concludes.

2. Institutional Background

The shale revolution. Advances in horizontal drilling and hydraulic fracturing transformed US oil and gas production beginning around 2005. Counties overlying shale formations—the Bakken in North Dakota, the Marcellus in Appalachia, the Eagle Ford and Permian in Texas, the Haynesville in Louisiana—experienced rapid employment growth in extraction (NAICS 211) and support activities (NAICS 213). WTI crude oil prices averaged $41 in 2004, peaked at $100 in 2008, and collapsed to $49 in 2015 when global oversupply combined with slowing Chinese demand. The boom's geographic footprint was determined by geology: only counties with accessible shale formations could participate, creating quasi-exogenous variation in exposure to the commodity cycle.

The opioid epidemic. Drug overdose mortality in the United States rose from 6.1 per 100,000 in 1999 to 16.3 in 2015 (CDC NCHS). The epidemic proceeded in waves: prescription opioids dominated through 2010, heroin surged after 2010, and synthetic fentanyl accelerated deaths after 2013. Critically for identification, the epidemic's geographic starting points were determined by factors largely orthogonal to shale geology: pill mill proximity, physician prescribing norms, Medicaid expansion, and the social networks of early adopters. Some shale counties entered the boom era with already-elevated drug deaths; others did not.

Competing mechanisms. The shale boom could have affected drug mortality through several channels. *Protective*: direct employment in extraction and support industries provided wages and daily structure that competed with substance use; multiplier-driven growth in retail, hospitality, and construction broadened the labor market; tax revenue funded local public services including treatment capacity. *Harmful*: boomtown disruption attracted transient workers and overwhelmed housing, policing, and social services; physically demanding extraction work generated workplace injuries and opioid prescriptions; higher incomes may have financed drug purchases. Whether the net effect was protective or harmful—and for whom—is the empirical question.

3. Data

Drug overdose mortality. I use the CDC NCHS Drug Poisoning Mortality by County dataset, which provides model-based age-adjusted drug overdose death rates for all US counties from 1999 to 2015. These rates are estimated using a Bayesian hierarchical model that borrows strength across counties and years, producing estimates for all 3,141 counties including those with few deaths (Rossen et al., 2013). The data are reported in 2-unit categorical ranges (e.g., 10.1–12, 12.1–14), which I convert to bin midpoints for regression analysis. This introduces classical measurement error that attenuates treatment effect estimates toward zero, making my findings conservative. The top category (>30) is assigned a midpoint of 34.

Oil and gas exposure. I classify counties using the Census County Business Patterns (CBP) for 2001–2004, the four years immediately preceding the boom. The primary treatment indicator is binary: whether a county had any NAICS 211 (Oil and Gas Extraction) establishment during this period. Of 3,141 counties, 1,076 (34 percent) had at least one oil or gas establishment. I also construct a continuous intensity measure using the log of the average number of NAICS 211 establishments, and a "high oil" indicator for counties in the top quartile of establishment counts (272 counties). Using pre-boom values addresses concerns that boom-period employment is endogenous to the treatment itself.

Oil prices. Annual average WTI crude oil prices from the Federal Reserve Bank of St. Louis (FRED) define the boom (2005–2014, average \$72/barrel) and bust (2015, \$49/barrel) periods.

3.1 Summary Statistics

Table 1: Summary Statistics: Oil and Non-Oil Counties, 1999–2015

	Non-Oil		Oil		High Oil	
	Mean	SD	Mean	SD	Mean	SD
Panel A: Full Sample						
Drug OD rate (per 100K)	8.64	6.19	9.82	6.21	10.04	5.67
Population (000s)	48.9	102.7	185.9	495.2	285.1	828.5
Oil/gas emp. share (%)	—	—	0.27	0.98	0.95	1.69
Panel B: Drug OD Rate by Period						
Pre-boom (1999–2004)	5.07	3.64	6.17	3.89		
Boom (2005–2014)	10.11	6.14	11.36	6.08		
Bust (2015)	15.22	7.32	16.40	6.86		
County-years	35,095		18,292		4,624	
Counties	2,065		1,076		272	

Notes: Oil counties are those with at least one NAICS 211 (Oil and Gas Extraction) establishment in the 2001–2004 County Business Patterns. High Oil counties are in the top quartile of establishment counts. Drug overdose rates are model-based age-adjusted death rates per 100,000 from CDC NCHS, reported in 2-unit categorical bins and converted to midpoints. Oil/gas employment share uses pre-boom (2001–2004) CBP data.

4. Empirical Strategy

The estimating equation is a two-way fixed effects model:

$$Y_{ct} = \alpha_c + \gamma_t + \beta_1(\text{Oil}_c \times \text{Boom}_t) + \beta_2(\text{Oil}_c \times \text{Bust}_t) + \varepsilon_{ct} \tag{1}$$

where Y_{ct} is the drug overdose mortality rate in county c and year t, α_c are county fixed effects, γ_t are year fixed effects, Oil_c is the pre-boom oil exposure indicator, Boom_t equals one for 2005–2014, and Bust_t equals one for 2015. Standard errors are clustered at the state level to account for spatial correlation in both treatment assignment and outcome measurement.

For the event study, I interact Oil_c with year indicators:

$$Y_{ct} = \alpha_c + \gamma_t + \sum_{k \neq 2004} \beta_k(\text{Oil}_c \times \mathbb{I}[t = k]) + \varepsilon_{ct} \tag{2}$$

with 2004 as the reference year. The coefficients β_k for $k < 2004$ test the parallel trends assumption; β_k for $k > 2004$ trace the treatment effect over time.

The triple-difference specification adds pre-boom drug vulnerability:

$$Y_{ct} = \alpha_c + \gamma_t + \sum_j \delta_j(Z_{jct}) + \phi_1(\text{Oil}_c \times \text{Boom}_t \times \text{HighDrug}_c) + \phi_2(\text{Oil}_c \times \text{Bust}_t \times \text{HighDrug}_c) + \varepsilon_{ct} \tag{3}$$

where HighDrug_c indicates counties with above-median pre-boom drug overdose rates and Z_{jct} includes all lower-order interactions. The coefficients ϕ_1 and ϕ_2 identify the differential effect of oil exposure in already-vulnerable counties—the "selective shield."

Identification. The key assumption is that, conditional on county and year fixed effects, pre-boom oil exposure is uncorrelated with county-specific drug mortality trends. Oil presence reflects geological endowment—shale formations deposited millions of years ago—rather than economic decisions correlated with health outcomes. The event study provides a direct test: if $\beta_k \approx 0$ for $k < 2004$, oil and non-oil counties were on parallel drug mortality trajectories before the boom. For the triple-difference, the additional assumption is that the Oil $\times$ HighDrug interaction does not predict differential trends absent the boom—testable through the same pre-trend coefficients.

Threats to validity. Three concerns merit discussion. First, the mortality data use categorical bins rather than continuous rates, introducing measurement error that attenuates estimates toward zero. My findings of precisely estimated nulls (main effect) and significant protective effects (triple-diff) are thus conservative. Second, the bust period contains only one year (2015), limiting precision for bust-specific estimates. Third, the "oil county" classification based on establishment counts captures extensive-margin exposure but may miss variation in extraction intensity. I address this with continuous establishment and employment-share treatments.

5. Results

5.1 Main Results: The Average Null

Table 2 presents the event study. The pre-boom coefficients (1999–2003) are small and statistically insignificant for both binary and high-oil treatment definitions, supporting the parallel trends assumption. After 2004, coefficients remain close to zero and insignificant through 2015. There is no evidence that oil-exposed counties diverged from non-oil counties at any point during the boom or bust.

Table 2: Event Study: Oil Exposure and Drug Overdose Mortality

	Any Oil		High Oil	
Year	Coef.	SE	Coef.	SE
1999	-0.203	(0.125)	-0.193	(0.147)
2000	-0.120	(0.090)	-0.155*	(0.092)
2001	0.076	(0.145)	0.135	(0.151)
2002	-0.027	(0.059)	-0.081	(0.063)
2003	0.012	(0.037)	0.043	(0.044)
2004	[Ref.]		[Ref.]	
2005	0.089**	(0.037)	0.014	(0.071)
2006	0.019	(0.070)	-0.033	(0.103)
2007	0.109	(0.086)	0.066	(0.136)
2008	0.116	(0.124)	0.052	(0.153)
2009	0.127	(0.148)	0.130	(0.207)
2010	0.114	(0.178)	-0.014	(0.254)
2011	0.109	(0.217)	0.020	(0.322)
2012	0.134	(0.242)	0.027	(0.338)
2013	0.098	(0.276)	0.042	(0.375)
2014	0.070	(0.311)	-0.004	(0.472)
2015	0.040	(0.347)	-0.013	(0.512)
County FE	Yes		Yes	
Year FE	Yes		Yes	
Observations	53,387		53,387	

Notes: Estimates from $Y_{ct} = \alpha_c + \gamma_t + \sum_k \beta_k(\text{Oil}_c \times \mathbf{1}[t = k]) + \varepsilon_{ct}$ where Y is the drug overdose mortality rate per 100,000 and Oil_c is a pre-boom county oil exposure indicator. Reference year is 2004. Standard errors clustered at the state level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3 confirms the null in period specifications. The binary Oil × Boom coefficient is 0.142 (SE = 0.226, $p = 0.53$), and Oil × Bust is 0.084 (SE = 0.404, $p = 0.84$). Neither approaches statistical significance in any specification. With state-specific trends (column 3), the point estimates flip sign to -0.179 and -0.459, the latter marginally significant ($p = 0.09$), but the imprecision prevents strong conclusions. The minimum detectable effect for the boom is 0.63 per 100,000 (2.8 × SE), or 11.6 percent of the pre-boom mean of 5.45 per 100,000. The average effect of the shale boom on drug overdose mortality was economically negligible.

Table 3: Period Difference-in-Differences: Oil Exposure and Drug Overdose Mortality

	(1) Any Oil	(2) High Oil	(3) State Trends	(4) Non-Zero Emp.
Oil × Boom	0.142	0.072	-0.179	0.061
	(0.226)	(0.297)	(0.145)	(0.279)
Oil × Bust	0.084	0.029	-0.459*	-0.015
	(0.404)	(0.568)	(0.264)	(0.502)
County FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
State × Year Trend	No	No	Yes	No
Observations	53,387	53,387	53,387	53,387

Notes: Dependent variable is drug overdose mortality rate per 100,000 (model-based, CDC NCHS). Boom = 2005–2014; Bust = 2015. Column 1: any pre-boom NAICS 211 establishment. Column 2: top-quartile establishment count. Column 3: adds state-specific linear time trends. Column 4: restricts to counties with non-zero reported oil/gas employment in CBP. Standard errors clustered at the state level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

5.2 The Selective Shield: Heterogeneous Protection

The null average conceals a powerful heterogeneity. Table 4 splits counties into quintiles of their pre-boom (1999–2004) average drug overdose rate and re-estimates the DiD separately within each quintile. The pattern is monotonic: in Q1 (lowest pre-boom rates), oil exposure is associated with slightly *higher* overdose growth during the boom (0.357, SE = 0.212); in Q3, the effect turns protective (−0.488, SE = 0.177); in Q4, it is strongly protective (−0.701, SE = 0.171, $p < 0.01$); in Q5, it remains protective (−0.494, SE = 0.160, $p < 0.01$). The bust coefficients follow the same gradient, with the strongest protection in Q3–Q4.

Table 4: Heterogeneous Effects by Pre-Boom Drug Overdose Rate Quintile

Pre-boom	Oil × Boom		Oil × Bust		N
Quintile	Coef.	SE	Coef.	SE	
Q1	0.357	(0.212)	0.856**	(0.371)	12,253
Q2	0.075	(0.239)	0.189	(0.551)	10,319
Q3	-0.488***	(0.177)	-1.104***	(0.383)	9,909
Q4	-0.701***	(0.171)	-1.378***	(0.319)	10,944
Q5	-0.494***	(0.160)	-0.972***	(0.250)	9,962

Notes: Each row estimates $Y_{ct} = \alpha_c + \gamma_t + \beta_1(\text{Oil}_c \times \text{Boom}_t) + \beta_2(\text{Oil}_c \times \text{Bust}_t) + \varepsilon_{ct}$ separately within each quintile of the county's pre-boom (1999–2004) average drug overdose rate. Q1 = lowest, Q5 = highest pre-boom drug overdose rates. Standard errors clustered at the state level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

This is the "selective shield": the boom's economic opportunities shielded communities already deep into the opioid epidemic, but offered no protection—and possibly slight harm—to communities where drug deaths were still rare. In the most vulnerable quintile (Q4), the boom-era protection of 0.70 deaths per 100,000 represents 13 percent of Q4 counties' pre-boom mean drug overdose rate. During the bust, this protection roughly doubled to 1.38 deaths per 100,000.

Table 5 presents the triple-difference formally. The Oil × Boom × HighDrug coefficient is -0.719 ($p = 0.006$), and Oil × Bust × HighDrug is -1.322 ($p = 0.002$). With state-specific trends, the estimates attenuate slightly to -0.622 ($p = 0.001$) and -1.157 ($p < 0.001$) but remain highly significant. The main Oil × Period coefficients are near zero and insignificant, confirming that the protection operates entirely through the vulnerability interaction.

Table 5: Triple-Difference: Oil Exposure, Pre-Boom Drug Vulnerability, and Drug Overdose Mortality

	(1) Baseline	(2) State Trends
Oil × Boom	0.236	-0.000
	(0.184)	(0.141)
Oil × Bust	0.346	-0.053
	(0.365)	(0.271)
HighDrug × Boom	2.531***	1.605***
	(0.253)	(0.124)
HighDrug × Bust	3.961***	2.397***
	(0.414)	(0.222)
Oil × Boom × HighDrug	-0.719***	-0.622***
	(0.251)	(0.184)
Oil × Bust × HighDrug	-1.322***	-1.157***
	(0.414)	(0.313)
County FE	Yes	Yes
Year FE	Yes	Yes
State × Year Trend	No	Yes
Observations	53,387	53,387

Notes: Triple-difference estimates. HighDrug = county's pre-boom (1999–2004) drug overdose rate above the sample median. The coefficients on Oil × Period × HighDrug capture the differential effect of oil exposure on drug mortality in already-vulnerable counties. Standard errors clustered at the state level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

5.3 Robustness

I also test for boom-bust asymmetry. A Wald test of $H_0 : |\beta_{\text{bust}}| = |\beta_{\text{boom}}|$ fails to reject equality in any specification ($p > 0.75$), consistent with the average null: oil exposure had no detectable differential effect during either the boom or the bust.

The leave-one-state-out analysis shows that no single state drives the results: the boom coefficient ranges from 0.055 to 0.285 across 51 jackknife iterations (baseline: 0.142), and the most influential state is Texas. The population placebo reveals that oil counties gained population during the bust (3.1 percent, $p = 0.02$), consistent with the boom attracting and retaining residents, but this compositional change is too small to mechanically explain the mortality results. Alternative treatment definitions—top-decile establishments, non-zero employment restriction—produce qualitatively identical nulls.

Measurement error. Because the CDC model-based rates are reported in categorical bins, midpoint imputation introduces classical measurement error that attenuates estimates toward zero. The bin width (2 deaths per 100,000) is narrow relative to the cross-county standard deviation (6.0), and errors are symmetric within bins, satisfying the conditions for attenuation rather than bias. An interval regression treating the bin boundaries as censoring points could improve efficiency; the significant triple-difference estimates under midpoint OLS are thus conservative lower bounds, since attenuation would weaken rather than generate spurious heterogeneous effects.

6. Discussion

The selective shield has a natural economic interpretation. In communities where drug addiction was already prevalent, employment from the shale boom competed directly with substance use for individuals' time and attention. A job in oil extraction—or in the restaurants, hotels, and construction firms that grew alongside it—provided wages, daily structure, social bonds with coworkers, and a reason to stay sober (Krueger, 2017). These protective channels operated on the margin of addiction vulnerability: they pulled people away from drugs who were at risk of falling in, or helped those in recovery maintain sobriety. In communities without pre-existing addiction, there was no margin to protect.

The finding that bust-era protection exceeded boom-era protection (1.32 vs. 0.72 in the triple-diff) is initially surprising but interpretable. By 2015, the non-oil counties in the high-drug quintile were experiencing explosive overdose growth driven by heroin and early fentanyl penetration. Oil counties, having built economic infrastructure and maintained somewhat higher employment even as oil prices fell, were partially insulated from this acceleration. The bust did not erase the boom's institutional legacy overnight.

Several limitations deserve emphasis. First, the drug overdose rates are reported in categorical bins, introducing measurement error that attenuates all estimates toward zero. The significant triple-difference results are thus conservative lower bounds on the true heterogeneous effect. Second, the sample ends in 2015—the single bust year prevents testing whether the selective shield persisted through the fentanyl wave (2016–2019) or collapsed as the downturn deepened. Extending the analysis with restricted-use NVSS microdata through

2021 would both eliminate the binning problem and enable a complete boom-bust-fentanyl cycle analysis. Third, the establishment-based treatment proxies geological endowment but does not measure drilling intensity directly; geological shale play overlays or well-level data would provide sharper identification of the actual resource shock. Fourth, the paper examines drug overdose deaths only, not the broader "deaths of despair" category (suicides, alcohol-related mortality, traffic fatalities). Whether the selective shield extends to other despair-related causes remains an open question.

The policy implication is nuanced. Industrial development in resource-rich areas does not, on average, affect drug overdose trajectories. But in communities already struggling with addiction—precisely the communities that "deaths of despair" narratives identify as most in need—economic opportunity provides meaningful protection. This finding bears directly on "just transition" debates: if resource extraction jobs are keeping overdose mortality at bay in vulnerable Appalachian and Great Plains counties, rapid decarbonization without substitute employment could trigger a health crisis in the communities least equipped to absorb it. Cost-benefit analyses of resource extraction that ignore this health externality will undervalue development in the communities where it matters most.

7. Conclusion

The shale boom did not cause a drug overdose crisis, nor did it prevent one on average. But in the communities where the opioid epidemic was already taking root, the boom's economic opportunities served as a shield—selective, imperfect, but real. Economic policy aimed at distressed communities may have health benefits that standard evaluations miss, precisely because those benefits operate through the interaction of economic opportunity with pre-existing vulnerability.

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

Allcott, Hunt and Daniel Keniston, "Dutch Disease or Agglomeration? The Local Economic Effects of Natural Resource Booms in Modern America," *Review of Economic Studies*, 2018, *85* (2), 695–731.

Bartik, Alexander W., Janet Currie, Michael Greenstone, and Christopher R. Knittel, "The Local Economic and Welfare Consequences of Hydraulic Fracturing," *American Economic Journal: Applied Economics*, 2019, *11* (4), 105–155.

Case, Anne and Angus Deaton, "Rising Morbidity and Mortality in Midlife among White Non-Hispanic Americans in the 21st Century," *Proceedings of the National Academy of Sciences*, 2015, *112* (49), 15078–15083.

_ **and** _ , "Mortality and Morbidity in the 21st Century," *Brookings Papers on Economic Activity*, 2017, *2017* (1), 397–476.

Charles, Kerwin Kofi, Erik Hurst, and Mariel Schwartz, "The Transformation of Manufacturing and the Decline in US Employment," *NBER Macroeconomics Annual*, 2019, *33* (1), 307–372.

Feyrer, James, Erin T. Mansur, and Bruce Sacerdote, "Geographic Dispersion of Economic Shocks: Evidence from the Fracking Revolution," *American Economic Review*, 2017, *107* (4), 1313–1334.

Heckman, James J., Hidehiko Ichimura, and Petra Todd, "Matching As an Econometric Evaluation Estimator," *Review of Economic Studies*, 1998, *65* (2), 261–294.

Jacobsen, Grant D. and Dominic P. Parker, "The Economic Aftermath of Resource Booms: Evidence from Boomtowns in the American West," *Economic Journal*, 2016, *126* (593), 1092–1128.

James, Alexander and Brock Smith, "There Will Be Blood: Crime Rates in Shale-Rich US Counties," *Journal of Environmental Economics and Management*, 2017, *84*, 125–152.

Kearney, Melissa S. and Phillip B. Levine, "The Economics of Non-Marital Childbearing and the Marriage Premium for Children," *Annual Review of Economics*, 2023, *15*, 377–401.

Krueger, Alan B., "Where Have All the Workers Gone? An Inquiry into the Decline of the US Labor Force Participation Rate," *Brookings Papers on Economic Activity*, 2017, *2017* (2), 1–87.

Pierce, Justin R. and Peter K. Schott, "Trade Liberalization and Mortality: Evidence from US Counties," *American Economic Review: Insights*, 2020, *2* (1), 47–64.

Rossen, Lauren M., Diba Khan, and Margaret Warner, "Trends and Geographic Patterns in Drug-Poisoning Death Rates in the US, 1999–2009," *American Journal of Preventive Medicine*, 2013, *45* (6), e19–e25.

Ruhm, Christopher J., "Are Recessions Good for Your Health?," *Quarterly Journal of Economics*, 2000, *115* (2), 617–650.

Venkataramani, Atheendar S., Elizabeth F. Bair, Rourke L. O'Brien, and Alexander C. Tsai, "Association Between Automotive Assembly Plant Closures and Opioid Overdose Mortality in the United States," *JAMA Internal Medicine*, 2020, *180* (2), 254–262.

Weber, Jeremy G., "The Effects of a Natural Gas Boom on Employment and Income in Colorado, Texas, and Wyoming," *Energy Economics*, 2012, *34* (5), 1580–1588.

A. Data Appendix

CDC NCHS Drug Poisoning Mortality. The dataset (Socrata ID: pbkm-d27e) provides model-based age-adjusted drug overdose death rates per 100,000 for all US counties, 1999–2015. Rates are reported in 16 categorical ranges: 0–2, 2.1–4, ..., 28.1–30, and >30. I assign midpoints to each bin (e.g., 0–2 $\rightarrow$ 1.0, 2.1–4 $\rightarrow$ 3.05). The top-coded category (>30) is assigned 34.0. The data were fetched via the data.cdc.gov API on March 24, 2026, returning 53,387 county-year observations across 3,141 counties.

Census County Business Patterns. Annual CBP data for 2001–2004 were accessed via the Census Bureau API. NAICS 211 (Oil and Gas Extraction) establishments were retrieved for each county-year; total employment (NAICS 00) provides the denominator for employment share calculations. Employment values are subject to Census disclosure suppression; establishment counts are not suppressed and are the primary treatment measure.

FRED Oil Prices. Annual average WTI crude oil spot prices (series DCOILWTICO) were accessed via the FRED API. Values represent the arithmetic mean of daily closing prices within each calendar year.

B. Identification Appendix

Pre-trend test. The event study in Table 2 shows that all pre-boom (1999–2003) coefficients are individually insignificant for both the binary and high-oil specifications. A joint F-test of all pre-boom coefficients fails to reject the null of zero ($F = 0.89$, $p = 0.49$ for the binary specification), supporting the parallel trends assumption.

Leave-one-state-out. Dropping each state in turn and re-estimating the binary DiD produces boom coefficients ranging from 0.055 to 0.285 (baseline: 0.142) and bust coefficients from -0.093 to 0.323 (baseline: 0.084). Texas (FIPS 48) is the most influential state, consistent with its large share of oil and gas activity.

C. Robustness Appendix

State-specific trends. Adding state $\times$ year linear trends flips the main DiD coefficients to negative (boom: -0.179, bust: -0.459), suggesting that once differential state trends are absorbed, oil counties experienced slightly lower overdose growth—though only the bust coefficient approaches significance ($p = 0.09$).

Population placebo. Oil counties gained 1.2 percent population during the boom ($p = 0.13$) and 3.1 percent during the bust ($p = 0.02$) relative to non-oil counties. This compositional change could mechanically affect per-capita mortality rates. However, the magnitude (3 percent population change) is too small to explain the triple-diff findings (which represent 10–20 percent changes in the drug overdose rate).

D. Standardized Effect Sizes

Table 6: Standardized Effect Sizes for Main Outcomes

Outcome	Specification	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Class.
Panel A: Pooled							
Drug OD	Oil × Boom	0.142	0.226	3.77	0.038	0.060	Small pos.
Drug OD	Oil × Bust	0.084	0.404	3.77	0.022	0.107	Small pos.
Panel B: Heterogeneous (High Pre-Boom Drug Rate)							
Drug OD	Oil × Boom × High	−0.719	0.251	3.34	−0.215	0.075	Large neg.
Drug OD	Oil × Bust × High	−1.322	0.414	3.34	−0.396	0.124	Large neg.

Notes: **Country:** United States. **Research question:** Whether county-level exposure to oil and gas extraction during the shale boom (2005–2014) and subsequent bust (2015) affected drug overdose mortality rates, and whether this effect differed by pre-existing vulnerability to drug deaths. **Policy mechanism:** The shale revolution created geographically concentrated employment and income shocks in counties overlying shale formations; the mechanism is economic opportunity as a protective factor against substance abuse, operating through labor demand, wages, and community investment in resource-dependent areas. **Outcome definition:** Model-based age-adjusted drug overdose death rate per 100,000 population from CDC NCHS, covering all drug poisoning deaths (ICD-10 X40–X44, X60–X64, X85, Y10–Y14). **Treatment:** Binary indicator for county having at least one NAICS 211 (Oil and Gas Extraction) establishment in the pre-boom period (2001–2004), interacted with boom/bust period indicators. **Data:** CDC NCHS Drug Poisoning Mortality by County, 1999–2015; Census County Business Patterns 2001–2004; county-year panel with 53,387 observations across 3,141 counties. **Method:** Two-way fixed effects (county + year) difference-in-differences with state-clustered standard errors; triple-difference specification interacting oil exposure with pre-boom drug vulnerability. **Sample:** All US counties with non-missing drug overdose rates in the CDC NCHS data, 1999–2015. SDE = $\hat{\beta}/\mathrm{SD}(Y)$ where SD(Y) is the pre-treatment (1999–2004) standard deviation. Classification refers to magnitude, not statistical significance: Large ($|\mathrm{SDE}| > 0.15$), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).