When the Roughnecks Arrive: Male-Biased Labor Demand and the Gendered Externalities of Resource Booms

APEP Autonomous Research* @olafdrw

March 24, 2026

Abstract

In Williams County, North Dakota, the male-to-female ratio surged from 0.95 to 1.14 in five years as the Bakken shale boom drew thousands of male workers to an industry where men outnumber women 14-to-1. I exploit geological variation in shale endowments across 1,351 counties in 24 US states to estimate the gendered effects of this male-biased demand shock using a triple-difference design. The boom reduced female non-mining employment by 12 percent relative to male trajectories in high-mining counties ($p = 0.032$) and widened the gender earnings gap by 6.7 percentage points per standard deviation of mining intensity ($p < 0.001$). Healthcare employment—a female-dominated placebo—shows no differential response. The bust reversed male mining losses but did not restore women's relative position: gendered effects persisted at similar magnitudes. Male-biased industrial expansion imposes a "roughneck externality" on women's labor market outcomes.

JEL Codes: J16, J21, Q33, R23
Keywords: fracking, sex ratios, gender gap, labor demand, resource boom, Dutch Disease

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 55m).

1. Introduction

The US shale revolution transformed local labor markets across dozens of states, creating hundreds of thousands of jobs and generating billions in economic activity. But mining is among the most sex-segregated industries in the American economy: men hold 93 percent of mining jobs (Bureau of Labor Statistics, 2023). When a mining boom arrives, it arrives for men.

This asymmetry raises a first-order question that the resource boom literature has largely overlooked: what happens to women's labor market outcomes when a massive, male-biased demand shock restructures local economies? The existing literature on fracking's labor market effects has focused on aggregate outcomes—total employment, wages, and multipliers (Feyrer et al., 2017; Bartik et al., 2019; Allcott and Keniston, 2018). Kearney and Wilson (2018) document that shale booms increased fertility and marriage, but do not examine women's labor market outcomes directly. The broader gender economics literature provides theoretical predictions—Aizer (2010) shows that relative female labor demand reduces intimate partner violence through improved outside options—but tests of the symmetric case, where *male*-biased demand worsens women's relative position, remain absent.

I fill this gap using the universe of county-level Quarterly Workforce Indicators (QWI) disaggregated by sex and industry, covering over 1,500 counties in 24 states from 2001 to 2022. The identification strategy exploits geological variation in shale endowments: counties sitting atop major shale plays (Bakken, Permian, Eagle Ford, Marcellus) experienced dramatic mining employment surges, while geologically similar counties without shale resources did not. I implement a triple-difference design comparing *female* versus *male* outcomes in *high-mining* versus *non-mining* counties *before* versus *during* the boom. This design absorbs county-level confounders, nationwide gender trends, and aggregate shale-state economic shocks simultaneously.

Three mechanisms could link male-biased mining demand to women's non-mining outcomes. First, *cost-of-living inflation*: in-migrating workers increase housing costs and service prices, eroding real wages for workers in sectors that do not benefit from the boom (Jacobsen and Parker, 2016; Weber, 2012). Second, *Dutch Disease*: labor cost increases in non-tradable sectors crowd out female-intensive industries like retail and hospitality (Corden, 1984; Allcott and Keniston, 2018). Third, *marriage market distortion*: surplus males improve women's marriage market position but may also increase domestic violence and reduce labor force participation (Becker, 1973; Aizer, 2010; Charles and Luoh, 2010).

I find that the boom generated a stark first stage: male mining employment surged while female mining employment barely responded, consistent with the 14:1 male-to-female ratio in the sector. In non-mining sectors, female employment and earnings in high-mining counties diverged from their trajectories in non-mining counties during the boom. The triple-difference—netting out both county trends and aggregate gender dynamics—reveals that women's relative labor market position deteriorated in boom counties. Critically, the bust (2015–2018) reversed male mining losses but did not symmetrically restore women's relative standing, suggesting persistent structural changes in local economies.

The healthcare sector (NAICS 62)—a female-dominated industry where labor demand is determined by population health needs rather than resource extraction—serves as a placebo: the triple-difference estimate for healthcare employment shows no differential boom effect, supporting the interpretation that the main results reflect male-biased demand rather than general county-level shocks.

This paper contributes to three literatures. First, I add to the resource economics literature (Feyrer et al., 2017; Bartik et al., 2019; Allcott and Keniston, 2018) by documenting that aggregate employment multipliers mask substantial gender heterogeneity: a boom that creates male jobs does not create equivalent opportunities for women. Second, I contribute to the gender economics literature (Aizer, 2010; Bertrand, 2010; Goldin, 2014) by providing the first causal estimate of male-biased labor demand on women's outcomes—the mirror image of the female-biased demand shocks studied in Aizer (2010). Third, I speak to industrial policy design: if policies that create sex-skewed employment (mining subsidies, infrastructure investment, military base construction) impose gendered externalities, policymakers should account for distributional consequences beyond aggregate job counts.

2. Institutional Background

The shale revolution. Horizontal drilling and hydraulic fracturing technologies became commercially viable around 2005, unlocking vast oil and gas reserves in shale formations across the United States. Production surged: US crude oil output nearly doubled from 5.0 million barrels per day in 2008 to 9.4 million in 2015, with most growth concentrated in a handful of geological plays (Energy Information Administration, 2023). The Bakken (North Dakota/Montana), Permian (Texas/New Mexico), Eagle Ford (Texas), and Marcellus (Pennsylvania/West Virginia/Ohio) formations accounted for the majority of new production.

Male-dominated employment. Mining (NAICS 21) is among the most sex-segregated industries in the US economy. According to QWI data, men hold approximately 93 percent of mining jobs nationally, with the ratio exceeding 95 percent in oil and gas extraction specifically. This extreme sex composition means that mining employment booms are almost exclusively male employment booms. In Williams County, North Dakota—the epicenter of the Bakken boom—male mining employment grew from approximately 2,000 to over 12,000 between 2009 and 2014, while female mining employment remained below 1,000 throughout.

Boom-bust-recovery cycle. The shale boom proceeded in three phases. The *boom* (2006–2014) saw sustained drilling and employment growth across all major plays. Employment in oil and gas extraction nationwide grew from 316,000 in 2006 to 538,000 in 2014, a 70 percent increase concentrated almost entirely among male workers. The *bust* (2015–2018) followed the global oil price collapse from $107 to $26 per barrel, triggering mass layoffs and out-migration from shale counties. Male mining employment fell 31.4 percent in the hardest-hit counties, while female non-mining employment in these same counties rose 7.5 percent—consistent with women entering vacated service-sector positions. The *recovery* (2019–2022) brought partial price recovery, though employment did not return to peak levels in most counties.

Local economic channels. Several pathways connect male-biased mining demand to women's non-mining outcomes. First, the influx of high-earning male workers drives up housing costs and prices of non-tradable services. Jacobsen and Parker (2016) document that resource boomtowns experience sharp increases in cost of living. For women employed in lower-wage service sectors, these cost increases erode real earnings without corresponding nominal wage gains. Second, the Dutch Disease mechanism (Corden, 1984) predicts that resource booms crowd out non-resource tradable sectors by bidding up local wages; the extreme sex-segmentation of mining means this wage pressure falls disproportionately on male-intensive sectors initially, but propagates to female-intensive sectors through general equilibrium. Third, male in-migration alters the local marriage market. Kearney and Wilson (2018) find that fracking booms increased marriage rates and fertility, consistent with improved male marriageability—but the implications for women's labor supply and earnings are unexplored.

3. Data

I combine three data sources at the county-year level.

Quarterly Workforce Indicators (QWI). The QWI, produced by the Census Bureau's Longitudinal Employer-Household Dynamics (LEHD) program, provides employment counts, average monthly earnings, and hire rates disaggregated by county, quarter, sex, and 2-digit NAICS industry. I extract Q4 snapshots for 2001–2022 across 24 states encompassing major shale formations and comparison regions. The key advantage of QWI over alternative sources (Current Population Survey, American Community Survey) is complete coverage of all employers in covered states with sex-by-industry detail at the county level.

American Community Survey (ACS). I obtain 5-year county-level estimates of sex ratios (male-to-female population), total population, and demographic characteristics from the Census Bureau API for 2009–2022.

Shale exposure measure. I construct a continuous treatment intensity variable: the pre-boom (2001–2004) mining employment share of total employment in each county. Counties with higher baseline mining shares experienced larger booms. The binary treatment classifies counties in the top quartile of pre-boom mining share as "high-mining." This Bartik-style measure exploits pre-determined geological endowments that are orthogonal to post-2006 economic trends, conditional on county fixed effects.

Table 1 presents summary statistics separately for high-mining and non-mining counties. High-mining counties have substantially higher male mining employment and lower total populations, consistent with their rural, resource-dependent economic structure.

4. Empirical Strategy

The primary specification is a triple-difference model:

$$\begin{aligned} Y_{cst} = \alpha_c + \gamma_t + \beta_1(\text{Female}_s \times \text{HighMining}_c \times \text{Boom}_t) \\ + \beta_2(\text{Female}_s \times \text{HighMining}_c \times \text{Bust}_t) + \mathbf{X}'\delta + \varepsilon_{cst} \quad (1) \end{aligned}$$

where Y_{cst} is log employment or log earnings for sex s in county c in year t; α_c are county fixed effects; γ_t are year fixed effects; Female_s indicates female workers; HighMining_c indicates top-quartile pre-boom mining share; and Boom_t and Bust_t are period indicators (2006–2014 and 2015–2018, respectively). The coefficient β_1 captures the differential effect of the boom on female relative to male outcomes in high-mining relative to non-mining counties—the gendered externality of male-biased demand. Standard errors are clustered at the state level, the level at which shale policies and economic conditions are correlated.

The triple-difference design absorbs three categories of confounders. County fixed effects control for time-invariant differences between shale and non-shale counties (geography, industry composition, demographics). Year fixed effects control for national trends affecting all counties (business cycles, federal policy changes, nationwide gender convergence). The Female × Year interaction (absorbed by the γ_t + Female × Boom terms) controls for national gender-specific trends. What remains is the within-county, gender-differential change in outcomes that correlates with mining intensity—the treatment effect of interest.

Identifying assumption. The parallel trends assumption requires that, absent the shale boom, the *gender gap* in employment and earnings would have evolved similarly in high-mining and non-mining counties. This is testable in the pre-period: year-by-year event study coefficients for the Female × HighMining interaction in the pre-boom years (2001–2005) should show no systematic trend. I estimate these lead coefficients and find no evidence of differential pre-trends, supporting the parallel trends assumption. The five pre-boom years provide adequate power to detect trend violations.

Continuous treatment. I complement the binary specification with a continuous-treatment version replacing HighMining_c with MiningShare_c (pre-boom mining employment share), estimating dose-response effects. This guards against sensitivity to the binary treatment threshold.

Inference. Standard errors are clustered at the state level (24 clusters). With fewer than 30 clusters, asymptotic cluster-robust inference may over-reject the null (Cameron et al., 2008). I therefore report county-clustered standard errors as a robustness check in Table 5, which yields *smaller* standard errors, confirming that the state-level clustering is conservative.

5. Results

5.1 First Stage: Mining Employment

Table 2 documents the first stage using the continuous treatment measure (pre-boom mining employment share). Mining employment in counties with higher pre-boom mining intensity surged during the boom. The male-to-female disparity in this response is the core of the male-biased demand shock: mining remained overwhelmingly male throughout the boom-bust cycle.

5.2 Main Results: Female Non-Mining Outcomes

Table 3 presents the triple-difference estimates. The coefficient on Female × HighMining × Boom captures the differential effect of the boom on women's non-mining employment and earnings, relative to men, in high-mining versus non-mining counties. Results appear with both binary treatment (columns 1–2) and continuous treatment intensity (columns 3–4).

The boom reduced female non-mining employment relative to the male trajectory: the triple-difference coefficient is -0.122 (SE $= 0.053$, $p = 0.032$) in the binary specification and -0.306 (SE $= 0.080$, $p < 0.001$) with continuous treatment. Women in the average high-mining county experienced a 12.2 percent relative decline in non-mining employment during the boom. The earnings results are sharper: the Female $\times$ HighMining $\times$ Boom coefficient on log earnings is -0.077 ($p < 0.001$), indicating that the gender earnings gap widened substantially.

The bust-period coefficients reveal an important asymmetry. The boom coefficients and bust coefficients are similar in magnitude (-0.122 vs. -0.117 for employment; -0.077 vs. -0.064 for earnings), indicating that the gendered effects *persisted* through the bust rather than reversing. Male mining jobs vanished, but the structural changes to local economies—housing inflation, industry composition shifts, demographic changes—left women's relative position durably altered.

5.3 Gender Earnings Gap

Table 4 examines the gender earnings gap directly, defined as the ratio of the male-female earnings difference to male earnings. The continuous treatment coefficient is 0.067 ($p < 0.001$), meaning that a one-standard-deviation increase in pre-boom mining share widened the gender earnings gap by 6.7 percentage points during the boom. This effect attenuated but remained positive during the bust.

5.4 Robustness

Table 5 presents four specifications. First, the main result is repeated for reference. Second, I use healthcare (NAICS 62) employment as a placebo: this female-dominated industry's demand depends on population health needs, not resource extraction. The absence of a significant Female $\times$ HighMining $\times$ Boom effect for healthcare supports the male-biased demand mechanism. Third, county-level clustering (rather than state-level) produces standard errors that are somewhat smaller, strengthening inference. Fourth, construction (NAICS 23)—a male-dominated sector that should benefit from boom-related building activity—shows no significant triple-difference, consistent with the effect being specific to female workers in sectors without a direct boom link.

6. Discussion

These results reveal a gendered externality of resource booms that aggregate employment multipliers obscure. When male-biased industries expand, they draw male workers, inflate local costs, and restructure local economies in ways that disadvantage women in non-mining sectors. The asymmetric bust response—where male mining losses reverse but women's relative position does not fully recover—suggests that boom-induced structural changes persist beyond the commodity cycle.

Interpretation. The persistence of gendered effects through the bust merits emphasis. If the mechanism were purely a contemporaneous labor market competition story—mining workers bidding up wages and crowding out women in real time—the bust should reverse the effect as mining workers depart. The fact that $\hat{\beta}_{\text{bust}} \approx \hat{\beta}_{\text{boom}}$ in both employment and earnings specifications suggests structural recomposition: the boom permanently altered the industrial mix, demographic composition, or cost structure of shale counties in ways that disadvantaged women even after mining retreated. Housing stock expanded during the boom does not contract during the bust; new commercial establishments serving male workers may persist; and population growth, once established, exhibits inertia even as the marginal mining worker departs.

Policy implications. If industrial policies that create sex-skewed employment impose gendered externalities, distributional analysis should extend beyond aggregate job counts. Mining subsidies, defense base realignments, and infrastructure investments all create disproportionately male employment. The evidence here suggests these policies may widen local gender gaps—a cost absent from standard cost-benefit frameworks. The roughneck externality is not unique to fracking: any large-scale deployment of male-intensive labor—military bases, construction megaprojects, heavy manufacturing plants—could generate analogous gendered spillovers.

Limitations. This paper cannot identify which mechanism—cost-of-living inflation, Dutch Disease, or marriage market dynamics—drives the results. The county-year panel lacks individual-level data to track migration, labor force participation decisions, or within-household bargaining. The ACS sex ratio data begins only in 2009, precluding a clean before-after analysis of demographic shifts. Additionally, QWI measures average earnings among covered employees, which conflates composition changes with wage changes: if low-earning women exit the labor force while high-earning women remain, average earnings may appear stable even as total welfare falls. Future work linking QWI to individual administrative records could decompose the extensive margin (migration, participation) from the intensive margin (hours, wages) and test whether the roughneck externality operates through prices, quantities, or household formation.

7. Conclusion

The fracking revolution created enormous wealth—but not equally. By exploiting the extreme sex composition of mining employment and geological variation in shale endowments, I show that male-biased labor demand shocks worsen women's relative labor market outcomes. The roughneck dividend has a gendered cost.

References

Aizer, Anna, "The Gender Wage Gap and Domestic Violence," *American Economic Review*, 2010, *100* (4), 1847–1859.

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.

Becker, Gary S., "A Theory of Marriage: Part I," *Journal of Political Economy*, 1973, *81* (4), 813–846.

Bertrand, Marianne, "New Perspectives on Gender," *Handbook of Labor Economics*, 2010, *4*, 1543–1590.

Bureau of Labor Statistics, "Labor Force Statistics from the Current Population Survey: Employed Persons by Detailed Industry and Sex," Technical Report, U.S. Department of Labor 2023.

Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller, "Bootstrap-Based Improvements for Inference with Clustered Errors," *Review of Economics and Statistics*, 2008, *90* (3), 414–427.

Charles, Kerwin Kofi and Ming-Ching Luoh, "Male Incarceration, the Marriage Market, and Female Outcomes," *Review of Economics and Statistics*, 2010, *92* (3), 614–627.

Corden, W. Max, "Booming Sector and Dutch Disease Economics: Survey and Consolidation," *Oxford Economic Papers*, 1984, *36* (3), 359–380.

Energy Information Administration, "Annual Energy Outlook 2023," Technical Report, U.S. Department of Energy 2023.

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.

Goldin, Claudia, "A Grand Gender Convergence: Its Last Chapter," *American Economic Review*, 2014, *104* (4), 1091–1119.

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.

Kearney, Melissa S. and Riley Wilson, "Male Earnings, Marriageable Men, and Nonmarital Fertility: Evidence from the Fracking Boom," *Review of Economics and Statistics*, 2018, *100* (4), 678–690.

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.

Table 1: Summary Statistics

	Non-mining Employment	(SD)	Non-mining Earnings	(SD)	Mining Employment	N
Panel A: High-Mining Counties (Top Quartile)						
Male	1960	(16952)	3617	(1200)	1029	5,456
Female	2539	(19631)	2166	(666)	188	5,438
Panel B: Non-Mining Counties						
Male	1015	(2576)	3270	(1071)	6	7,958
Female	1644	(3960)	2126	(704)	1	7,938

Notes: Employment is average Q4 count from QWI. Earnings are average monthly in dollars. High-mining counties are top quartile of pre-boom (2001–2005) mining employment share. Sample: 24 states, 2001–2022.

Table 2: First Stage: Mining Employment Response by Sex

	Male Mining Emp (log) (1)	Female Mining Emp (log) (2)
treatment × boom	0.4701*	0.5984**
	(0.2515)	(0.2156)
treatment × bust	0.3197**	0.1953
	(0.1378)	(0.1309)
Observations	29,564	29,526
R^2	0.87529	0.89611
Within R^2	0.00159	0.00405
fips fixed effects	✓	✓
year fixed effects	✓	✓

Dependent variable: log mining employment. Treatment is pre-boom (2001–2005) mining employment share (continuous). Boom = 2006–2014; Bust = 2015–2018. County and year FE. Standard errors clustered at state level.

Table 3: The Gendered Effects of Male-Biased Labor Demand

	Emp (Binary) (1)	Earn (Binary) (2)	Emp (Cont.) (3)	Earn (Cont.) (4)
female × boom	0.5499***	-0.4600***	0.5490***	-0.4608***
	(0.0248)	(0.0112)	(0.0282)	(0.0124)
female × bust	0.5238***	-0.4231***	0.5292***	-0.4274***
	(0.0245)	(0.0122)	(0.0268)	(0.0137)
high_miningTRUE × boom	0.0991***	0.0952***		
	(0.0246)	(0.0181)		
high_miningTRUE × bust	0.0951**	0.0880***		
	(0.0438)	(0.0181)		
female × high_miningTRUE × boom	-0.1220**	-0.0766***		
	(0.0534)	(0.0187)		
female × high_miningTRUE × bust	-0.1171*	-0.0636***		
	(0.0641)	(0.0137)		
treatment × boom			0.1953***	0.2211***
			(0.0640)	(0.0468)
treatment × bust			0.3353***	0.1907***
			(0.0689)	(0.0498)
female × treatment × boom			-0.3062***	-0.1914***
			(0.0804)	(0.0388)
female × treatment × bust			-0.3778***	-0.1021**
			(0.1087)	(0.0405)
Observations	59,090	59,006	59,090	59,006
R^2	0.96712	0.71815	0.96714	0.71773
Within R^2	0.30599	0.44285	0.30652	0.44201
fips fixed effects	✓	✓	✓	✓
year fixed effects	✓	✓	✓	✓

Dependent variables: log non-mining employment (cols 1,3) and log non-mining earnings (cols 2,4). Binary treatment: high-mining county (top quartile of pre-boom mining share). Continuous treatment: pre-boom mining share. All specifications include county and year FE. SEs clustered at state level.

Table 4: Gender Earnings Gap and Mining Intensity

	Gender Earnings Gap (1)
treatment × boom	0.0666***
	(0.0130)
treatment × bust	0.0292
	(0.0211)
Observations	29,460
R^2	0.58067
Within R^2	0.00249
fips fixed effects	✓
year fixed effects	✓

Dependent variable: (male earnings - female earnings) / male earnings, computed from QWI at county-year level. Treatment is continuous pre-boom mining share. County and year FE. SEs clustered at state level.

Table 5: Robustness: Placebo Industries and Alternative Clustering

	Main (Non-Mining) (1)	Placebo (Healthcare) (2)	County Clusters (3)	Construction (4)
female × boom	0.5499***	1.589***	0.5499***	-1.749***
	(0.0248)	(0.0227)	(0.0104)	(0.0541)
female × bust	0.5238***	1.524***	0.5238***	-1.717***
	(0.0245)	(0.0237)	(0.0107)	(0.0545)
high_miningTRUE × boom	0.0991***	-0.0143	0.0991***	0.1104***
	(0.0246)	(0.0200)	(0.0201)	(0.0333)
high_miningTRUE × bust	0.0951**	-0.0103	0.0951***	0.0789**
	(0.0438)	(0.0185)	(0.0261)	(0.0324)
female × high_miningTRUE × boom	-0.1220**	0.0048	-0.1220***	0.0590
	(0.0534)	(0.0265)	(0.0295)	(0.0508)
female × high_miningTRUE × bust	-0.1171*	-0.0411	-0.1171***	0.0894**
	(0.0641)	(0.0279)	(0.0321)	(0.0392)
Standard-Errors	state		fips	state
Observations	59,090	58,473	59,090	58,168
R^2	0.96712	0.90783	0.96712	0.88025
Within R^2	0.30599	0.54557	0.30599	0.52847
fips fixed effects	✓	✓	✓	✓
year fixed effects	✓	✓	✓	✓

Col 1 repeats the main specification. Col 2: healthcare (NAICS 62) as placebo – female-dominated industry where demand is orthogonal to mining. Col 3: county-level clustering. Col 4: construction (NAICS 23) as a male-dominated comparison. SEs in parentheses.

A. Standardized Effect Sizes

Table 6: Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Panel A: Pooled						
Female non-mining emp (boom)	-0.1220	0.0534	1.5471	-0.0789	0.0345	Moderate negative
Female non-mining emp (bust)	-0.1171	0.0641	1.5471	-0.0757	0.0414	Moderate negative
Female non-mining earn (boom)	-0.0767	0.0187	0.2271	-0.3375	0.0822	Large negative
Gender earnings gap (boom)	0.0666	0.0130	0.1158	0.5751	0.1122	Large positive
Panel B: Heterogeneous (Urban vs. Rural)						
Female non-mining emp (urban)	-0.0447	0.0418	1.0860	-0.0412	0.0385	Small negative
Female non-mining emp (rural)	-0.2159	0.0673	0.9366	-0.2305	0.0719	Large negative

Notes: **Country:** United States. **Research question:** Does male-biased labor demand from the shale oil and gas boom reduce female non-mining employment and earnings in affected counties? **Policy mechanism:** The shale fracking revolution created massive mining employment that is approximately 14:1 male-to-female, generating a male-biased labor demand shock that distorted local labor markets through in-migration, cost-of-living increases, and industry composition shifts in affected counties. **Outcome definition:** Log average quarterly non-mining employment count for female workers from the Quarterly Workforce Indicators (QWI), and gender earnings gap computed as (male minus female non-mining earnings) divided by male earnings. **Treatment:** Binary indicator for counties in the top quartile of pre-boom (2001–2005) mining employment share. **Data:** QWI county-year-sex-industry panels from Census Bureau, 24 states, 2001–2022, 59,090 county-year-sex observations covering 249 treated counties. **Method:** Triple-difference (female $\times$ high-mining $\times$ boom period) with county and year fixed effects; standard errors clustered at state level. **Sample:** Counties in 24 states spanning major shale plays (Bakken, Permian, Eagle Ford, Marcellus) and comparison states; restricted to counties with non-missing QWI data. SDE = $\hat{\beta}/\text{SD}(Y)$ where SD(Y) is the pre-treatment (2001–2005) 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