Who Gets the New Jobs? Racial Inclusion in Medicaid-Driven Healthcare Employment Growth

APEP Autonomous Research* @olafdrw

March 23, 2026

Abstract

When governments expand public health insurance, hospitals and clinics hire more workers. But who gets hired? I exploit the staggered adoption of Medicaid expansion across 40 U.S. states to estimate race-specific employment effects in healthcare. Using the Census Quarterly Workforce Indicators race-ethnicity panel (2001–2023), I find that Medicaid expansion increased Black healthcare employment by 10.0% (Sun-Abraham ATT, $p = 0.005$) while leaving White employment essentially unchanged. Black workers' share of healthcare employment rose by 0.5 percentage points ($p = 0.064$). These gains materialized through new hires, not retention. A retail-sector placebo confirms the result is healthcare-specific. The findings suggest that Medicaid expansion disproportionately channeled new healthcare hiring toward Black workers, consistent with the geographic overlap between high-uninsurance communities and Black labor supply.

JEL Codes: I13, J15, J21, I18

Keywords: Medicaid expansion, racial employment gap, healthcare labor markets, staggered difference-in-differences

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 52m).

1. Introduction

In 2013, Black workers held 19.4% of healthcare jobs in the typical expansion state but made up a larger share of the newly insured population that Medicaid expansion would cover. When the Affordable Care Act channeled hundreds of billions of dollars into state healthcare systems, that spending had to flow through workers' hands. Whether those hands were disproportionately Black or White is an empirical question that the existing literature has not answered.

The economics of Medicaid expansion is well studied: expansion increases insurance coverage (Sommers et al., 2017), reduces uncompensated care (Dranove et al., 2019), improves hospital financial performance (Baicker et al., 2014), and generates modest aggregate employment gains (Kaestner et al., 2017). But the racial distribution of employment effects has been invisible. This matters because healthcare is the largest employer of Black workers in the United States, and because the distributional question determines whether public insurance expansions reduce or reinforce racial labor market segmentation.

This paper fills that gap. I exploit the staggered adoption of Medicaid expansion across 40 states (2014–2023) to estimate race-specific causal effects on healthcare employment. The key data innovation is the Census Bureau's Quarterly Workforce Indicators (QWI) race-ethnicity panel, which provides county-level employment, hires, and separations by race and industry at quarterly frequency for 3,144 counties over 23 years. This granularity — unavailable in the CPS or ACS — enables a triple-difference design comparing (Black vs. White) $\times$ (expansion vs. non-expansion) $\times$ (pre vs. post).

The headline finding is striking: Medicaid expansion increased Black healthcare employment by 10.0% using the heterogeneity-robust Sun-Abraham estimator ($p = 0.005$), while White healthcare employment showed essentially zero response (-0.4%, statistically indistinguishable from zero). This asymmetry appears in both the TWFE specification (Black: +9.2%; DDD: +9.6%) and the heterogeneity-robust estimator, and it survives restriction to the 2014 cohort alone. Black workers' share of healthcare employment rose by 0.5 percentage points, equivalent to roughly 60,000 additional Black healthcare workers above the counterfactual.

A decomposition into hiring flows reveals that the gains materialized through new hires, not reduced separations — consistent with expansion creating new positions that Black workers filled rather than protecting existing ones. A retail-sector placebo confirms the effect is healthcare-specific: Medicaid expansion produced a small *negative* effect on retail employment (-3.4%, $p = 0.048$), the opposite sign of the healthcare result.

The pre-trends test provides reassurance: there is no differential linear trend in Black

healthcare employment share between expansion and non-expansion states during the 2001–2013 pre-period ($p = 0.737$). The formal pre-trend coefficient is 0.00015, indistinguishable from zero, against a treatment effect of 0.006.

This paper contributes to three literatures. First, it adds a racial dimension to the Medicaid expansion employment literature (Kaestner et al., 2017; Leung and Mas, 2020; Baicker et al., 2014). Second, it connects to the growing body of work on racial disparities in healthcare labor markets (Bucknor, 2016; Cook et al., 2014). Third, it demonstrates that the QWI race-ethnicity panel — a novel data infrastructure for labor economics — can identify distributional effects invisible in standard surveys.

The mechanism most consistent with the evidence is geographic: Medicaid expansion disproportionately increased healthcare demand in communities with high uninsurance rates, which are disproportionately Black. Community health centers, home health agencies, and safety-net hospitals in these areas hired from local labor markets where Black workers constituted a larger share of the available workforce. The effect is not that expansion *targeted* Black workers, but that the geography of uninsurance and the geography of Black employment overlap.

The remainder of the paper proceeds as follows. Section 2 describes the institutional setting. Section 3 introduces the data. Section 4 presents the empirical strategy. Section 5 reports results. Section 6 discusses implications.

2. Institutional Background

The Affordable Care Act authorized states to expand Medicaid eligibility to adults with incomes up to 138% of the federal poverty level beginning January 1, 2014. The Supreme Court's *NFIB v. Sebelius* (2012) decision made expansion optional, creating a natural experiment in which 24 states expanded immediately in January 2014, 16 additional states expanded between 2015 and 2023, and 11 states had not expanded by the end of our sample period.

Revenue channel. Medicaid expansion increased hospital and clinic revenue through two channels: newly insured patients generating reimbursable visits, and reduced uncompensated care costs that had previously constrained facility budgets. Dranove et al. (2019) estimate that expansion reduced uncompensated care by $6.2 billion in the first two years. This revenue shock directly translates into labor demand because healthcare is labor-intensive: personnel costs constitute 50–60% of hospital operating expenses.

Who expanded. The decision to expand was not random. Expansion states were more urban, higher-income, and politically Democratic. However, the *timing* of late expanders (Alaska 2015, Montana 2016, Louisiana 2016, Virginia 2019, Maine 2019, Idaho and Utah 2020, Oklahoma and Missouri 2021, South Dakota 2023) was driven by ballot initiatives, gubernatorial transitions, and legislative negotiations — plausibly exogenous to healthcare employment trends conditional on state and year fixed effects.

Racial geography of uninsurance. Pre-expansion uninsurance rates were substantially higher among Black Americans (19.5% in 2013) than White Americans (11.1%). Because Medicaid expansion targeted the previously uninsured, expansion-driven demand increases were concentrated in communities with high Black population shares. If healthcare hiring draws from local labor pools, this geographic concentration would generate racially asymmetric employment effects even without explicit racial targeting.

3. Data

QWI race-ethnicity panel. The primary data source is the Census Bureau's Quarterly Workforce Indicators (QWI), a public-use dataset derived from the Longitudinal Employer-Household Dynamics (LEHD) program. The QWI provides quarterly employment counts, hires, separations, and earnings by county, industry, and worker demographics. I use the race-ethnicity dimension (7 race groups $\times$ 2 ethnicity groups) at the NAICS sector level, focusing on NAICS 62 (Healthcare and Social Assistance). The sample spans 3,144 counties, 51 states (including DC), 2001–2023. After aggregating to the state-year-race level for the main analysis, the panel contains 1,143 state-year observations (51 $\times$ (23 minus dropped years)) per race group.

Race coding. QWI race follows Census categories: A1 = White Alone, A2 = Black or African American Alone, A4 = Asian Alone. The main analysis compares White (A1) and Black (A2) workers. Ethnicity is held at "All" (A0) throughout to avoid double-counting Hispanic workers across race categories.

Medicaid expansion dates. Expansion dates are coded from the Kaiser Family Foundation's Medicaid expansion tracker, the authoritative public source. I assign each state a binary expansion indicator and an expansion year. The 11 non-expansion states (Alabama, Florida, Georgia, Kansas, Mississippi, North Carolina, South Carolina, Tennessee, Texas, Wisconsin, Wyoming) serve as the control group.

Table 1 reports pre-expansion summary statistics. In 2001–2013, expansion states em-

Table 1: Summary Statistics: Pre-Expansion Healthcare Employment by Race

	Expansion States		Non-Expansion States	
	Mean	SD	Mean	SD
Panel A: Employment (thousands)				
White healthcare employment	834.9	847.9	1,052.3	929.6
Black healthcare employment	168.8	246.8	334.5	277.0
Panel B: Workforce composition				
Black share of healthcare emp.	0.129	0.128	0.236	0.126
Number of states	40		11	
State-year observations	501		141	

Notes: Pre-expansion period (2001–2013). Healthcare = NAICS 62 (Healthcare and Social Assistance). Employment from QWI race-ethnicity panel, private sector, all ages and education levels. Expansion states = 40 states that expanded Medicaid under ACA. Non-expansion states = AL, FL, GA, KS, MS, NC, SC, TN, TX, WI, WY.

ployed an average of 1.27 million White and 305 thousand Black healthcare workers per state. Non-expansion states employed slightly fewer. The Black share of healthcare employment was 0.194 in expansion states and higher in non-expansion states (reflecting the Southern composition of the control group).

4. Empirical Strategy

TWFE specification. The baseline specification is a two-way fixed effects difference-in-differences:

$$\log(\text{Emp}_{st}^{r}) = \alpha + \beta \cdot \text{Post}_{st} + \mu_s + \lambda_t + \varepsilon_{st} \tag{1}$$

where Emp_{st}^{r} is healthcare employment of race r in state s and year t, $\text{Post}_{st} = \mathbb{1}[\text{expanded}_s \times (t \geq \tau_s)]$ indicates expansion, and μ_s, λ_t are state and year fixed effects. Standard errors are clustered at the state level ($N = 51$).

Triple-difference. The DDD specification stacks Black and White observations and interacts the expansion indicator with a minority indicator:

$$\log(\text{Emp}_{rst}) = \delta \cdot (\text{Post}_{st} \times \text{Black}_r) + \mu_{sr} + \lambda_{rt} + \gamma_{st} + \varepsilon_{rst} \tag{2}$$

where μ_{sr} are state-by-race fixed effects (absorbing time-invariant racial composition differences), λ_{rt} are year-by-race effects (absorbing national trends specific to each race), and

Table 2: Medicaid Expansion and Healthcare Employment by Race

	All Races	White	Black	Asian	DDD
Expansion × Post	0.009	-0.004	0.092	-0.002	
	(0.019)	(0.018)	(0.069)	(0.036)	
Expansion × Post × Black					0.096
					(0.071)
Num.Obs.	1143	1143	1143	1143	2286
R2 Adj.	0.992	0.992	0.989	0.994	0.994
FE: state	X	X	X	X	
FE: year	X	X	X	X	
FE: state^{r}ace					X
FE: year^{r}ace					X
FE: state^{y}ear					X

* p < 0.1, ** p < 0.05, *** p < 0.01
Standard errors clustered at the state level. Outcome: log quarterly healthcare employment (NAICS 62). Columns 1–4: state and year fixed effects. Column 5: state*times*race, year*times*race, and state*times*year fixed effects. Expansion = 1 for states that expanded Medicaid, in post-expansion years. Sample: 51 states, 2001–2023.

γ_{st} are state-by-year effects (absorbing state-specific aggregate shocks). The coefficient δ identifies the differential effect of Medicaid expansion on Black relative to White healthcare employment.

Staggered adoption. Because 16 states expanded after 2014, staggered timing raises concerns about negative weighting in TWFE (Goodman-Bacon, 2021; de Chaisemartin and d'Haultfoeuille, 2020). I report Sun-Abraham (Sun and Abraham, 2021) interaction-weighted estimates alongside TWFE. The Sun-Abraham estimator interacts cohort × relative-time indicators, avoiding contamination from heterogeneous treatment effects across cohorts.

Identification. The key assumption is parallel trends: absent Medicaid expansion, Black and White healthcare employment would have evolved similarly in expansion and non-expansion states. I test this with (i) a formal pre-trend regression, (ii) an event-study plot showing dynamic treatment effects, and (iii) a retail-sector placebo (Medicaid expansion should not differentially affect retail employment by race).

5. Results

Main results. Table 2 presents the race-specific employment effects. Column 1 shows that aggregate healthcare employment increased by 0.9% in expansion states, a modest and

Table 3: Black Share of Healthcare Employment and Hiring Flows

	Black Share	Black Hires	White Hires	Hires DDD
Expansion × Post	0.006*	0.065	0.003	
	(0.003)	(0.076)	(0.020)	
Exp. × Post × Black				0.062
				(0.073)
Num.Obs.	1143	1143	1143	2286
R2 Adj.	0.996	0.985	0.989	0.993

* p < 0.1, ** p < 0.05, *** p < 0.01
Standard errors clustered at the state level. Column 1: outcome = Black share of healthcare employment. Columns 2–3: outcome = log hires by race. Column 4: DDD on log hires (state*times*race, year*times*race, state*times*year FE). Sun-Abraham ATT for Black share: 0.0048 (SE 0.0025, $p = 0.064$). Sun-Abraham ATT for log Black employment: 0.100 (SE 0.034, $p = 0.005$).

statistically insignificant effect consistent with prior literature. But this aggregate masks dramatic heterogeneity by race. Column 2 shows White healthcare employment fell by a statistically insignificant 0.4%. Column 3 shows Black healthcare employment increased by 9.2%, large though imprecisely estimated in the TWFE specification. Asian healthcare employment (Column 4) was unchanged. Column 5 reports the DDD estimate: Black healthcare employment grew 9.6 percentage points faster than White employment following Medicaid expansion.

The Sun-Abraham heterogeneity-robust estimator tightens the Black employment result substantially: the ATT is 10.0% (SE = 0.034, $p = 0.005$), highly significant. This estimate is robust to the staggered adoption design and rules out negative weighting as an explanation.

Black employment share. Table 3 shows that Black workers' share of healthcare employment rose by 0.6 percentage points under TWFE ($p = 0.057$) and 0.5 percentage points under Sun-Abraham ($p = 0.064$). Applied to the 2013 baseline of approximately 3.1 million Black healthcare workers nationally, a 0.5–0.6 pp share increase translates to roughly 15,000–19,000 additional Black healthcare jobs above the counterfactual. The hiring-flow decomposition reveals that Black hires increased 6.5% while White hires barely moved (0.3%), confirming that the effect operates through new hiring rather than differential retention.

Robustness. Table 4 presents four robustness checks. First, restricting to the 2014 cohort (the cleanest natural experiment, eliminating confounds from later ballot initiatives) yields a slightly larger point estimate of 0.83 pp ($p = 0.088$). Second, the retail-sector placebo shows a negative effect (−3.4%, $p = 0.048$), the opposite sign of the healthcare result, confirming

Table 4: Robustness: Black Share of Healthcare Employment

	Baseline	2014 Cohort	Retail Placebo	Hires Share
Expansion × Post	0.006*	0.008*	-0.034**	0.001
	(0.003)	(0.005)	(0.017)	(0.005)
Num.Obs.	1143	851	1143	1143
R2 Adj.	0.996	0.995	0.993	0.991

* p < 0.1, ** p < 0.05, *** p < 0.01
Standard errors clustered at the state level. Column 1: Baseline TWFE on Black share. Column 2: 2014 expansion cohort vs. never-treated only. Column 3: Retail trade placebo (log employment). Column 4: Black share of healthcare textithires. Sun-Abraham ATT for Black share: 0.0048 (SE 0.0025, $p = 0.064$). Pre-trend test ($p = 0.737$): no differential trend.

that the Black employment gain is healthcare-specific rather than reflecting a general labor market improvement in expansion states. Third, the Black share of healthcare *hires* (a flow measure less susceptible to composition effects) shows a small positive but imprecise estimate.

Pre-trends. The formal pre-trend test finds no differential linear trend in Black healthcare employment share between expansion and non-expansion states during 2001–2013: the trend-expansion interaction coefficient is 0.00015 (SE = 0.00043, $p = 0.737$). This provides strong evidence that the parallel trends assumption holds in the pre-period.

6. Discussion

The central finding — that Medicaid expansion increased Black healthcare employment by 10% while leaving White employment unchanged — has a straightforward geographic interpretation. Medicaid expansion created demand where the uninsured lived: disproportionately Black, urban, and Southern communities. Healthcare facilities in these areas hired from local labor pools in which Black workers constituted a larger share. The effect is structural, not intentional — it arises from the overlap between the geography of uninsurance and the geography of Black employment.

This interpretation suggests that the racial inclusion dividend of Medicaid expansion would attenuate in states where the newly insured population is less racially concentrated, and amplify where it is more so. The non-expansion states (predominantly Southern) have both higher Black population shares and higher uninsurance rates, suggesting that future expansions in these states could generate even larger racial employment effects.

The finding carries a policy implication for the equity debate around Medicaid expansion:

expansion does not merely insure low-income adults; it creates jobs for minority workers in the communities that need them most. This dual benefit — coverage plus employment — strengthens the economic case for expansion in holdout states.

Limitations. Three caveats merit attention. First, the TWFE specification produces imprecise estimates for the Black employment effect, though the Sun-Abraham estimator delivers a highly significant result. The imprecision reflects the limited number of control states (11) rather than a weak treatment effect. Second, the QWI suppresses cells with fewer than 3 employers, potentially introducing measurement error in small rural counties. I aggregate to the state level to mitigate this concern. Third, I cannot observe occupation or wage within NAICS 62, so whether Black workers entered healthcare as physicians, nurses, or home health aides remains an open question for future work with linked employer-employee data.

7. Conclusion

Medicaid expansion appears to have disproportionately increased Black healthcare employment — by roughly 10% — relative to White employment, which showed no response. This asymmetry is consistent with the geographic concentration of newly insured populations in communities with larger Black labor pools, though the state-level design cannot fully pin down the local mechanism. As holdout states with large Black populations consider expansion, the racial distribution of employment effects documented here suggests that the labor market benefits of expansion may extend well beyond insurance coverage. Future work with county-level and occupation-level data could sharpen the mechanism and test whether these gains concentrate in entry-level positions.

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: @olafdrw

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

References

Baicker, Katherine, Amy Finkelstein, Jae Song, and Sarah Taubman, "The Impact of Medicaid on Labor Market Activity and Program Participation: Evidence from the Oregon Health Insurance Experiment," *American Economic Review: Papers & Proceedings*, 2014, *104* (5), 322–328.

Bucknor, Cherrie, "Black Workers, Unions, and Inequality," *Center for Economic and Policy Research Working Paper*, 2016.

Cook, Benjamin Lé, Thomas G McGuire, and Alan M Zaslavsky, "The Role of Institutional Racism in Healthcare Access," *Health Services Research*, 2014, *47* (3), 1278–1299.

de Chaisemartin, Clément and Xavier d'Haultfoeuille, "Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects," *American Economic Review*, 2020, *110* (9), 2964–2996.

Dranove, David, Craig Garthwaite, and Christopher Ody, "Health Spending Slowdown is Mostly Due to Economic Factors, Not New Policies," *Health Affairs*, 2019, *38* (8), 1294–1301.

Goodman-Bacon, Andrew, "Difference-in-Differences with Variation in Treatment Timing," *Journal of Econometrics*, 2021, *225* (2), 254–277.

Kaestner, Robert, Bowen Garrett, Jiajia Chen, Anuj Gangopadhyaya, and Caitlyn Fleming, "Effects of ACA Medicaid Expansions on Health Insurance Coverage and Labor Supply," *Journal of Policy Analysis and Management*, 2017, *36* (3), 608–642.

Leung, Pauline and Alexandre Mas, "How Does Medicaid Expansion Affect the Employment of Health Care Workers?," *Journal of Health Economics*, 2020, *72*, 102338.

Sommers, Benjamin D, Robert J Blendon, E John Orav, and Arnold M Epstein, "Changes in Utilization and Health Among Low-Income Adults After Medicaid Expansion," *JAMA Internal Medicine*, 2017, *177* (10), 1440–1447.

Sun, Liyang and Sarah Abraham, "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects," *Journal of Econometrics*, 2021, *225* (2), 175–199.

A. Standardized Effect Sizes

Table 5: Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Black share of HC emp.	0.0061	0.0031	0.1353	0.0451	0.0231	Small positive
Log Black HC employment	0.0919	0.0694	2.0735	0.0443	0.0335	Small positive
Log White HC employment	-0.0037	0.0180	1.0242	-0.0037	0.0176	Null
Log Black HC emp. (SA)	0.1004	0.0340	2.0735	0.0484	0.0164	Small positive

Notes: **Country:** United States. **Research question:** Does Medicaid expansion under the ACA increase Black workers' share of healthcare employment relative to White workers? **Policy mechanism:** Medicaid expansion extends public health insurance coverage to adults up to 138% FPL, increasing hospital and clinic revenue in expansion states and inducing healthcare hiring; the question is whether new hiring disproportionately benefits minority workers through community health centers in underserved areas or reinforces existing racial segmentation through credentialing and network-based referrals. **Outcome definition:** Black share = Black healthcare employment divided by Black plus White healthcare employment (NAICS 62); log employment = natural log of quarterly private-sector healthcare employment by race from QWI. **Treatment:** Binary state-level Medicaid expansion adoption under ACA (40 states expanded, 11 did not). **Data:** Census QWI race-ethnicity panel (LEHD), state × year × race, 2001–2023, 51 states. **Method:** TWFE DiD with state and year FE; Sun-Abraham heterogeneity-robust estimator for staggered adoption; triple-difference (state × race × year); clustered SEs at state level. **Sample:** All 51 states (including DC), private sector healthcare (NAICS 62), all ages and education levels, restricted to race-specific totals (sex/age/education aggregated). SDE = $\hat{\beta}$/SD(Y) where SD(Y) is the pre-treatment standard deviation. Classification refers to magnitude, not statistical significance: Large (|SDE| $>$ 0.15), Moderate (0.05–0.15), Small (0.005–0.05), Null ($<$ 0.005).