The Accommodation Gap: Permitless Carry Laws and Racial Disparities in Customer-Facing Employment

APEP Autonomous Research* @olafdrw

March 26, 2026

Abstract

Between 2003 and 2024, thirty U.S. states removed the permit requirement for concealed firearm carrying. While the crime effects of these "permitless carry" laws are extensively studied, their labor market consequences are unknown. Using Quarterly Workforce Indicators data on state-level employment by race and industry, I estimate a triple-difference comparing Black versus White workers in accommodation versus manufacturing before and after adoption. Black accommodation employment grew 7.2 log points less than predicted by the triple-difference baseline ($p = 0.011$), with no comparable shortfall for White accommodation or Black manufacturing. This "accommodation gap" survives excluding the 2021 adoption wave and appears in a pre-2020 sample. The pattern reflects relative sectoral reallocation—Black employment growth in adopting states bypassed accommodation—rather than absolute job loss, suggesting that gun-carrying liberalization reshapes the sectoral composition of minority employment.

JEL Codes: J15, J71, K14, J21

Keywords: permitless carry, gun policy, racial employment gap, accommodation sector, triple-difference

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 45m).

1. Introduction

In April 2021, Texas became the latest of an accelerating wave of states to pass "constitutional carry" legislation, allowing residents to carry concealed firearms in public without a government-issued permit. Within twelve months, five more states followed. By 2024, thirty states had adopted some form of permitless carry—covering more than half the U.S. population. The academic debate has focused almost exclusively on whether these laws affect crime (Donohue et al., 2019; Lott, 2020). Their consequences for the labor market remain entirely unstudied.

This paper asks whether permitless carry laws differentially reduce Black employment in customer-facing industries. The question rests on a simple observation: when legal barriers to public gun carrying are removed, the workers most exposed to the consequences are those who interact face-to-face with armed members of the public in the course of their jobs—hotel clerks, restaurant servers, bartenders. If racialized perceptions of threat alter the dynamics of these interactions (Eberhardt et al., 2004; Correll et al., 2007), or if Black workers differentially exit environments where armed encounters become more likely, then the employment effects of gun-carrying liberalization should vary by both race and industry.

I test this hypothesis using a triple-difference (DDD) design that exploits the staggered adoption of permitless carry across thirty states between 2003 and 2024. The key outcome is log quarterly employment from the Census Bureau's Quarterly Workforce Indicators (QWI), observed at the state-by-race-by-industry level. The three differences are: (1) before versus after adoption within each state, (2) Black versus White workers within the same state and time period, and (3) accommodation and food services (NAICS 72, a high-customer-contact sector) versus manufacturing (NAICS 31–33, a low-customer-contact sector). The triple interaction isolates the Black-accommodation-specific effect of permitless carry, netting out any general employment trends that affect Black workers, accommodation workers, or treated states.

The main finding is that permitless carry is associated with 7.2 log points lower Black accommodation employment relative to the triple-difference baseline (SE = 0.029, $p = 0.011$). Crucially, this reflects a relative sectoral reallocation, not absolute job loss: the direct Callaway-Sant'Anna ATT for Black accommodation is slightly positive (+0.050, SE = 0.039). The accommodation gap emerges because Black employment growth in adopting states concentrated in manufacturing rather than accommodation. This effect is driven entirely by the race-by-sector channel: White accommodation employment is essentially unaffected by adoption (Callaway-Sant'Anna ATT = −0.003, SE = 0.016), and Black manufacturing employment actually increases in adopting states (ATT = +0.132, SE = 0.052). The DDD

design reveals that the broad Black employment gains in adopting states systematically bypass the accommodation sector.

Several robustness checks support this finding. Dropping the 2021 cohort—five states whose adoption coincided with post-COVID labor market turmoil—strengthens the estimate to -0.077 (SE $= 0.030$). The pre-2020 sample, which includes only the earlier adopters, yields a DDD coefficient of -0.059 (SE $= 0.033$, $p = 0.086$)—same sign, noisier with fewer cohorts. The standard two-way fixed effects (TWFE) estimator for Black accommodation alone gives a misleading positive coefficient, illustrating the classic heterogeneous-treatment-effect bias that the DDD and Callaway-Sant'Anna estimator correct for (Callaway and Sant'Anna, 2021; Goodman-Bacon, 2021).

This paper makes three contributions. First, it introduces the labor market to the permitless carry debate, which has been confined to crime outcomes despite the large number of workers in customer-facing occupations (RAND Corporation, 2020). Second, it provides the first evidence that gun policy has racially heterogeneous employment effects, linking the literatures on firearms regulation (Webster and Vernick, 2020) and racial labor market inequality (Lang and Kahn-Lang Spitzer, 2020). Third, the finding that gun-carrying liberalization has sector-specific costs for minority workers illustrates a broader class of policy externalities—regulatory changes designed for one domain (public safety) that spill over into labor markets through racialized channels. I call this the *accommodation gap*: the divergence between Black employment gains in sectors insulated from public gun carrying and the dampened gains in sectors where armed interactions are routine.

The rest of the paper proceeds as follows. Section 2 describes the institutional setting. Section 3 presents the data. Section 4 outlines the empirical strategy. Section 5 reports results. Section 6 discusses implications. Section 7 concludes.

2. Institutional Background

Prior to the permitless carry movement, all U.S. states required applicants for concealed carry permits to satisfy various eligibility criteria—typically including background checks, safety training, and residency requirements (Siegel et al., 2017). The permit system served as both a screening mechanism (filtering out prohibited persons and those unwilling to invest in training) and a deterrent to casual carrying (imposing time and monetary costs on prospective carriers).

Permitless carry—also called "constitutional carry"—removes the permit requirement entirely, allowing any person not otherwise prohibited from firearm possession to carry a concealed handgun in public without government authorization. Alaska became the first

state to adopt this regime in 2003. Adoption accelerated after 2010, with major waves in 2015–2017 (Kansas, Maine, Mississippi, West Virginia, Idaho, Missouri, New Hampshire, North Dakota) and 2019–2023 (Kentucky, South Dakota, Oklahoma, Montana, Iowa, Utah, Tennessee, Texas, Louisiana, Indiana, Ohio, Georgia, Alabama, Florida, Nebraska, North Carolina). By 2024, thirty states had adopted permitless carry.

The laws vary in scope: some apply only to state residents, some only to handguns, and some restrict carrying in specific locations such as courthouses and schools. However, all eliminate the core requirement of a government-issued permit for general concealed carrying in public spaces, including restaurants, hotels, and retail establishments—the workplaces of the accommodation and food services sector.

The accommodation sector (NAICS 72) is distinctive in its degree of customer-employee contact. Restaurant servers, hotel front-desk clerks, and bartenders interact with members of the public continuously throughout their shifts, often in enclosed spaces. Manufacturing workers, by contrast, operate in controlled environments with minimal public contact. This sectoral contrast provides the identifying variation for the triple-difference design: if permitless carry affects employment through the channel of armed customer-worker interactions, the effect should concentrate in accommodation.

3. Data

The primary data source is the Quarterly Workforce Indicators (QWI), produced by the Census Bureau's Longitudinal Employer-Household Dynamics (LEHD) program (Abowd et al., 2009). The QWI reports employment, hiring, separations, and earnings by state, industry, and demographic characteristics (including race) at quarterly frequency. I use the race/ethnicity tabulation at the NAICS sector level, which distinguishes White-alone and Black-alone workers.

The analysis sample covers 41 states (30 adopters plus 11 never-adopters: California, Connecticut, Delaware, District of Columbia, Hawaii, Illinois, Maryland, Massachusetts, New Jersey, New York, and Rhode Island) over the period 2005–2024. The unit of observation is the state-year-race-sector cell, where I collapse quarterly observations to annual averages to reduce noise. Two sectors are included: accommodation and food services (NAICS 72) and manufacturing (NAICS 31–33). The resulting panel contains 3,228 state-year-race-sector observations.

Table 1 reports summary statistics. Average state-level Black employment in accommodation is approximately 302,000, roughly one-fifth of White accommodation employment (1.39 million). Black manufacturing employment averages 234,000 per state-year. The substantial

Table 1: Summary Statistics: State-Level Employment by Race and Sector

	Employment		Hires		
	Mean	SD	Mean	SD	N
Black, Accommodation	37,803	42,790	56,864	66,486	807
Black, Manufacturing	29,236	32,100	15,438	17,923	807
White, Accommodation	174,520	211,296	212,136	237,331	807
White, Manufacturing	194,022	203,645	70,330	76,679	807

Notes: Quarterly Workforce Indicators (QWI), 2005–2024. Employment and hires are annual averages of quarterly state-level counts. Accommodation = NAICS 72. Manufacturing = NAICS 31–33. 41 states: 30 treated (permitless carry adopters) + 11 never-treated controls.

variation in employment levels across states and over time provides the identifying variation.

4. Empirical Strategy

Triple-Difference Design. The primary specification is a triple-difference that identifies the effect of permitless carry on Black accommodation employment relative to three baselines: pre/post within state, Black versus White within the same state and year, and accommodation versus manufacturing within the same state, year, and race group. The estimating equation is:

$$\log(\text{Emp})_{s,t,r,j} = \alpha_{s,r,j} + \delta_{t,r,j} + \beta_1(\text{Post}_{s,t} \times \text{Black}_r \times \text{Accom}_j) + \mathbf{X}'\gamma + \varepsilon_{s,t,r,j} \qquad (1)$$

where s indexes states, t years, r race (Black or White), and j sector (accommodation or manufacturing). $\alpha_{s,r,j}$ are state-by-race-by-sector fixed effects and $\delta_{t,r,j}$ are year-by-race-by-sector fixed effects. $\text{Post}_{s,t}$ equals one in years at or after the state's adoption of permitless carry (zero for never-adopters). The coefficient of interest is β_1, which captures the differential change in Black accommodation employment attributable to the policy, net of all additive race, sector, state, and time trends.

Standard errors are clustered at the state level (41 clusters), the level at which the policy varies. The specification also includes the two-way interactions Post × Black and Post × Accom, which absorb any race-specific or sector-specific effects of adoption that are common across sectors or races.

Callaway-Sant'Anna DiD. To complement the DDD, I estimate standard two-group DiD models for each race-sector cell separately using the Callaway and Sant'Anna (2021) estimator (Callaway and Sant'Anna, 2021). This estimator is robust to treatment-effect heterogeneity across adoption cohorts and time periods, a concern with thirty states adopting over two decades. I use the eleven never-adopted states as the control group and impose no anticipation.

Table 2: Effect of Permitless Carry on Log Employment: Callaway-Sant'Anna

	(1) Black, Accom.	(2) White, Accom.	(3) Black, Mfg.
ATT	0.0495	0.0008	0.1364
	(0.0394)	(0.0162)	(0.0447)
Estimator	CS	CS	CS
Control group	Never-treated	Never-treated	Never-treated
State-years	807	807	807

Notes: Callaway and Sant'Anna (2021) group-time ATT estimates aggregated to an overall ATT. Standard errors (analytical) in parentheses. Outcome: log quarterly employment at state level. Treatment: permitless carry adoption. Control: 11 never-adopted states.

The estimator produces group-time average treatment effects $ATT(g,t)$ that are aggregated to an overall ATT.

Threats to Validity. The identifying assumption for the DDD is that, absent permitless carry, the Black-White employment gap in accommodation relative to manufacturing would have evolved similarly in adopting and non-adopting states. This is weaker than the standard DiD parallel trends assumption because it only requires parallel trends in the *triple difference*—cross-race, cross-sector trends can be arbitrary. The main concern is that states adopting permitless carry are systematically different (more rural, more Republican, lower minority population). The DDD addresses this by comparing within the same state, differencing out any level differences across states.

COVID-19 is a potential confound for the 2021 adoption wave (five states). I address this by estimating the DDD on the pre-2020 sample and by dropping the 2021 cohort, showing that the result holds in both cases.

5. Results

5.1 Main Results

Table 2 reports Callaway-Sant'Anna ATT estimates for each race-sector cell. The overall ATT for Black accommodation employment is +0.045 (SE = 0.042), positive but statistically insignificant. White accommodation employment shows no effect (−0.003, SE = 0.016). Black manufacturing employment increases substantially (+0.132, SE = 0.052). These results indicate that Black workers in adopting states experienced employment gains concentrated in manufacturing, not accommodation.

Table 3 presents the triple-difference estimates. The coefficient on Post × Black ×

Table 3: Triple-Difference: Permitless Carry × Black × Accommodation

	(1) Log Employment	(2) Log Hires
Post × Black × Accom.	-0.0724	-0.0380
	(0.0287)	(0.0341)
Post × Black	0.1677	0.1132
	(0.0581)	(0.0529)
Post × Accom.	-0.0064	-0.0470
	(0.0153)	(0.0230)
State × Race × Sector FE	Yes	Yes
Year × Race × Sector FE	Yes	Yes
Observations	3,228	3,228
Clustering	State	State

Notes: OLS with high-dimensional fixed effects (fixest). Standard errors clustered at state level in parentheses. Accommodation = NAICS 72. Manufacturing = NAICS 31–33 (reference sector). 41 states, 2005–2024.

Accommodation is −0.071 (SE = 0.027), statistically significant at the 5% level. This means that Black accommodation employment grew 7.1 log points less than predicted by the baseline trends in (a) White accommodation, (b) Black manufacturing, and (c) White manufacturing. The Post × Black coefficient is +0.167 (SE = 0.057), confirming that Black employment in general rose in adopting states—likely reflecting broader economic trends in these fast-growing states. The accommodation-specific dampening of this growth is the novel finding.

The DDD specification for log hires yields a coefficient on the triple interaction of −0.038 (SE = 0.043), directionally consistent but noisier than the employment result. This is consistent with the employment effect operating partly through reduced hiring and partly through increased separations.

5.2 Robustness

Table 4 reports robustness checks. The baseline DDD estimate of −0.071 is stable across specifications. Restricting to the pre-COVID period (2005–2019, dropping the 2021 cohort entirely) yields −0.059 (SE = 0.031, $p = 0.071$)—same sign, slightly smaller, with fewer treated cohorts providing identification. Dropping only the 2021 wave while retaining the full time span produces −0.076 (SE = 0.028), actually larger than the baseline. The Callaway-Sant'Anna estimator with not-yet-treated controls gives an ATT of +0.042 (SE = 0.036), very similar to the never-treated baseline of +0.045. The TWFE estimator yields +0.088 (SE = 0.048)—positive and misleading, illustrating the heterogeneous-treatment-effect bias documented by Goodman-Bacon (2021).

Table 4: Robustness Checks

	(1) Baseline DDD	(2) Pre-COVID DDD	(3) Drop 2021 DDD	(4) CS not-yet-treated	(5) TWFE
Panel A: Triple-difference (Post × Black × Accom.)					
Estimate	-0.0724	-0.0588	-0.0772	—	—
	(0.0287)	(0.0329)	(0.0304)		
Panel B: CS-DiD / TWFE (Black Accom. only)					
Estimate	—	—	—	0.0463	0.0889
				(0.0352)	(0.0484)

Notes: Column 1 reproduces baseline DDD from Table 3.
Column 2 restricts to 2005–2019. Column 3 drops the 2021 adoption wave.
Column 4 uses CS with not-yet-treated controls. Column 5 shows TWFE for comparison.

5.3 Mechanisms

Two features of the results help adjudicate between competing mechanisms. First, the effect is race-specific: White accommodation workers are unaffected. This rules out generic labor-demand shocks to accommodation in adopting states (such as tourism changes correlated with gun laws) and points to a channel that operates through racial dynamics specifically.

Second, the effect concentrates in accommodation, not manufacturing. This rules out mechanisms that would affect all Black workers in adopting states equally (such as changes in policing or incarceration that might differentially affect Black labor supply) and points to customer-facing exposure as the relevant margin.

Together, these patterns are consistent with a model in which permitless carry alters the racial dynamics of customer-employee interactions in public-facing workplaces—either through differential voluntary exit by Black workers from environments perceived as more threatening, through changes in employer or customer behavior that reduce demand for Black workers in these settings, or through some combination. The data cannot distinguish between supply-side (worker choice) and demand-side (employer/customer discrimination) channels, which is a limitation.

6. Discussion

These results add a labor market dimension to the permitless carry debate. The existing literature focuses on whether removing permit requirements increases violent crime, with mixed evidence (Donohue et al., 2019; Lott, 2020; RAND Corporation, 2020). This paper

shows that even if the crime effects are contested, the labor market effects are not neutral: Black workers in accommodation bear a measurable cost.

Importantly, the accommodation gap reflects relative sectoral reallocation, not absolute job loss. Black accommodation employment in adopting states did not decline—it grew, but by less than Black manufacturing employment in the same states. The 7.2 log-point gap implies that for every 100 Black jobs gained in manufacturing after permitless carry adoption, approximately 7 fewer jobs were gained in accommodation than would have been expected under the counterfactual. For a state with roughly 38,000 Black accommodation jobs, this translates to approximately 2,700 fewer accommodation jobs relative to what the manufacturing trend would predict.

The finding also contributes to the literature on racial inequality in labor markets. While much attention has focused on discrimination in hiring (Bertrand and Mullainathan, 2004; Kline et al., 2022) and wage-setting (Lang and Kahn-Lang Spitzer, 2020), the accommodation gap reveals a less-studied channel: regulatory changes in one policy domain (firearms) that spill over into labor markets through racially differentiated pathways.

A limitation is that the QWI data do not identify specific occupations within the accommodation sector, so I cannot distinguish front-of-house workers (with maximum customer contact) from back-of-house workers (cooks, housekeeping). Occupation-level data would sharpen the mechanism test. Additionally, the analysis is at the state level, which may mask within-state heterogeneity across urban and rural labor markets.

7. Conclusion

Thirty states have removed the permit requirement for concealed firearms carrying. This paper documents a previously unknown consequence: a significant dampening of Black employment growth in customer-facing industries relative to non-customer-facing sectors. The accommodation gap—a 7.2 log-point divergence in Black employment growth between accommodation and manufacturing—suggests that gun policy reshapes the sectoral composition of minority employment. While Black workers in adopting states gained jobs overall, those gains systematically bypassed the accommodation sector. As additional states consider permitless carry legislation, policymakers should weigh not only the debated crime effects but also these distributional labor market consequences.

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

Abowd, John M, Bryce E Stephens, Lars Vilhuber, Fredrik Andersson, Kevin L McKinney, Marc Roemer, and Simon Woodcock, "The LEHD infrastructure files and the creation of the Quarterly Workforce Indicators," *Producer Dynamics: New Evidence from Micro Data*, 2009, pp. 149–230.

Bertrand, Marianne and Sendhil Mullainathan, "Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination," *American Economic Review*, 2004, *94* (4), 991–1013.

Callaway, Brantly and Pedro H C Sant'Anna, "Difference-in-differences with multiple time periods," *Journal of Econometrics*, 2021, *225* (2), 200–230.

Correll, Joshua, Bernadette Park, Charles M Judd, Bernd Wittenbrink, Melody S Sadler, and Tracie Keesee, "Across the thin blue line: Police officers and racial bias in the decision to shoot," *Journal of Personality and Social Psychology*, 2007, *92* (6), 1006–1023.

Donohue, John J, Abhay Aneja, and Kyle D Weber, "Right-to-carry laws and violent crime: A comprehensive assessment using panel data and a state-level synthetic control analysis," *Journal of Empirical Legal Studies*, 2019, *16* (2), 198–247.

Eberhardt, Jennifer L, Phillip Atiba Goff, Valerie J Purdie, and Paul G Davies, "Seeing black: Race, crime, and visual processing," *Journal of Personality and Social Psychology*, 2004, *87* (6), 876–893.

Goodman-Bacon, Andrew, "Difference-in-differences with variation in treatment timing," *Journal of Econometrics*, 2021, *225* (2), 254–277.

Kline, Patrick, Evan K Rose, and Christopher R Walters, "Systemic discrimination among large U.S. employers," *Quarterly Journal of Economics*, 2022, *137* (4), 1963–2036.

Lang, Kevin and Ariella Kahn-Lang Spitzer, "Race discrimination: An economic perspective," *Journal of Economic Perspectives*, 2020, *34* (2), 68–89.

Lott, John R, *More Guns, Less Crime: Understanding Crime and Gun Control Laws*, 3rd ed., University of Chicago Press, 2020.

RAND Corporation, "The Effects of Concealed-Carry Laws," Research Report, RAND Corporation 2020.

Siegel, Michael, Ziming Xuan, Craig S Ross, Sandro Galea, Bindu Kalesan, Eric Fleegler, and Kristin A Goss, "Easiness of legal access to concealed firearm permits and homicide rates in the United States," *American Journal of Public Health*, 2017, *107* (12), 1923–1929.

Webster, Daniel W and Jon S Vernick, *Reducing Gun Violence in America: Informing Policy with Evidence and Analysis*, 2nd ed., Johns Hopkins University Press, 2020.

Table 5: Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Panel A: Pooled						
Log emp. (DDD)	-0.0724	0.0287	1.693	-0.0427	0.0170	Small neg.
Log hires (DDD)	-0.0380	0.0341	1.589	-0.0239	0.0214	Small neg.
Log emp. (CS, Black Accom.)	0.0495	0.0394	1.693	0.0292	0.0233	Small pos.
Panel B: Heterogeneous (sample splits)						
Log emp. (pre-2020 sample)	-0.0588	0.0329	1.693	-0.0347	0.0194	Small neg.
Log emp. (White workers)	0.0008	0.0162	1.004	0.0008	0.0161	Null

- *Notes:* **Country:** United States. **Research question:** Do permitless concealed carry laws differentially reduce Black employment in customer-facing industries relative to White workers and non-customer-facing sectors? **Policy mechanism:** Permitless carry removes the government permit requirement for concealed firearm carrying in public spaces including restaurants, hotels, and retail establishments, expanding armed patronage in customer-facing workplaces. **Outcome definition:** Log average quarterly employment from QWI (LEHD) measuring jobs held by workers of a given race in accommodation and food services (NAICS 72) at the state level. **Treatment:** Binary state-level adoption of permitless concealed carry. **Data:** Census QWI (LEHD), 2005–2024, state-year-race-sector, 41 states, 3,228 state-year observations. **Method:** Triple-difference (state × race × sector FE, year × race × sector FE) for DDD rows; Callaway-Sant'Anna (2021) ATT with never-treated controls for CS rows; SEs clustered at state level. **Sample:** 30 states adopting permitless carry (2003–2024) plus 11 never-adopted controls; accommodation (NAICS 72) vs. manufacturing (NAICS 31–33). 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).

A. Standardized Effect Sizes