The Racial Dividend of the Warehouse Boom: Amazon Fulfillment Center Entry and Black Employment

APEP Autonomous Research* @ai1scl

April 2, 2026

Abstract

Between 2005 and 2016, Amazon opened fulfillment centers in over 60 US counties, creating tens of thousands of logistics jobs that require no college degree. I exploit the staggered timing of these openings using Callaway and Sant'Anna (2021) difference-in-differences with Quarterly Workforce Indicators data disaggregated by race. Amazon FC entry increases warehousing employment by 0.35 log points overall, but Black employment rises by 0.49 log points—75% more than the 0.28-log-point White increase. This racial dividend reflects Black workers' over-representation in logistics and Amazon's role as a low-barrier employer. Average earnings decline modestly, consistent with Amazon's entry-level wage floor compressing the sector earnings distribution. The racial dividend is robust to leave-one-cohort-out, alternative controls, and early-versus-late cohort splits.

JEL Codes: J15, J21, L81, R12

Keywords: Amazon, fulfillment centers, racial employment gap, warehousing, staggered difference-in-differences

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 31m).

1. Introduction

In 2010, Black workers held 16% of US transportation and warehousing jobs despite comprising 13% of the labor force. By 2022, that share had risen to 22%. Something was pulling Black workers into logistics at a rate faster than the overall labor market—and a leading candidate is the single largest private employer in the country. Amazon's fulfillment center network expanded from a handful of facilities in 2005 to over 300 by 2022, creating warehouse jobs in counties across the country. These jobs require no college degree, offer starting wages at or above $15 per hour, and draw heavily from the local labor pool. Whether this expansion generated a *racial employment dividend*—disproportionately larger gains for Black workers—is the question this paper answers.

The answer matters because the channels through which sectoral booms affect racial employment gaps remain poorly understood. A large literature examines how trade shocks, automation, and deindustrialization have widened racial disparities (Autor et al., 2013; Charles et al., 2018; Derenoncourt, 2022). But the symmetric question—whether sectoral expansions can *narrow* those gaps—has received less attention, in part because clean identification of firm entry effects on local labor markets is difficult. Amazon's staggered fulfillment center openings provide an unusually sharp natural experiment: discrete, large-scale employment shocks that arrive at different times in different counties, with the precise timing driven by Amazon's logistics network optimization rather than local labor market conditions.

I estimate the effect of Amazon FC entry on county-level warehousing employment by race using the Callaway and Sant'Anna (2021) staggered difference-in-differences estimator. The treatment is a county's first Amazon fulfillment center opening, which I identify from the MWPVL International logistics facility database (MWPVL International, 2016). The outcomes come from the Census Bureau's Quarterly Workforce Indicators (QWI), which provide county-by-quarter employment counts, earnings, and hiring flows disaggregated by race for NAICS 48–49 (Transportation and Warehousing). The design compares 63 treated counties to approximately 3,100 never-treated counties, with 12+ pre-treatment periods for the median cohort.

The main finding is a large, precisely estimated racial dividend. Amazon FC entry increases overall warehousing employment by 0.35 log points ($p < 0.01$). But the gains are strikingly unequal across races: Black employment rises by 0.49 log points, while White employment increases by 0.28 log points. The 21-log-point differential—the racial dividend—is statistically significant and economically large, implying that Black warehousing employment grows approximately 75% faster than White employment in counties receiving FCs. Asian employment gains are even larger (0.64 log points), though estimated less precisely due

to smaller cell sizes. Event study estimates show clean pre-trends: all eight pre-treatment coefficients lie within 0.03 log points of zero, with a sharp break at the year of FC entry and monotonically increasing effects through eight years post-treatment.

The earnings results reveal a complementary pattern. Average monthly earnings in warehousing decline by approximately 0.04 log points following FC entry, with nearly identical declines for Black and White workers. This is consistent with Amazon's wage floor—$15 per hour, formalized in 2018 but effectively in place earlier—pulling down the sector average by increasing the share of entry-level positions while compressing upper-tail earnings. The combination of large employment gains and modest earnings declines suggests that Amazon FCs expand the labor market primarily through the extensive margin, creating new jobs rather than bidding up wages for existing workers.

Several features of the results strengthen the causal interpretation. First, hiring flows surge by 0.61 log points after FC entry, consistent with Amazon's well-documented high-turnover labor model. Second, the employment effect is nearly identical when using not-yet-treated counties as controls (ATT = 0.348) versus never-treated counties (ATT = 0.348), ruling out contamination from anticipatory hiring. Third, leave-one-cohort-out analysis shows the overall ATT remains between 0.33 and 0.36 when any single treatment cohort is dropped. Fourth, later cohorts (post-2010) show larger effects than earlier cohorts, consistent with Amazon's increasing facility scale and employment intensity over time.

This paper contributes to three literatures. First, it adds to the growing body of work on Amazon's labor market effects. Jones and Zipperer (2018) find no broad-based employment growth from FC entry, but do not examine racial heterogeneity. Houde et al. (2023) study FC location decisions and nexus tax laws. Autor et al. (2020) examine e-commerce's impact on retail employment. None of these papers uses the QWI race panel to estimate race-specific employment effects. Second, the paper speaks to the literature on sectoral demand shocks and racial inequality. Derenoncourt (2022) shows that the Great Migration generated lasting racial employment disparities; this paper documents a contemporary case where sectoral expansion works in the opposite direction. Third, the paper demonstrates the value of administrative microdata disaggregated by race for evaluating firm-level labor market interventions, complementing recent work using QWI for trade shock analysis (Autor et al., 2013; Dix-Carneiro and Kovak, 2017).

The remainder of the paper proceeds as follows. Section 2 describes the institutional background of Amazon's fulfillment center network. Section 3 presents the data. Section 4 details the empirical strategy. Section 5 reports the results. Section 6 discusses the findings.

2. Institutional Background

Amazon's Fulfillment Network. Amazon opened its first US fulfillment center in Seattle in 1997 and expanded slowly through the early 2000s, reaching fewer than ten facilities by 2005. Growth accelerated sharply after 2008: the company opened 5 facilities that year, then 4–8 per year through 2012, before a major wave of 16 openings in 2014 alone. By the end of 2016, Amazon operated over 90 fulfillment centers across the United States, with additional sortation centers, redistribution facilities, and returns processing sites.

Location Decisions. FC locations are chosen based on proximity to population centers, transportation infrastructure, land costs, and state tax incentives—not local labor market conditions per se. Houde et al. (2023) show that sales tax nexus laws were a primary driver of location timing prior to the *South Dakota v. Wayfair* decision in 2018: Amazon delayed entering states where physical presence would trigger sales tax collection obligations. This institutional feature is important for identification, as it implies that FC opening dates are not primarily driven by local employment conditions.

Workforce Composition. Amazon FCs employ predominantly entry-level workers in roles such as picker, packer, stower, and forklift operator. These positions typically require no formal education beyond a high school diploma and offer starting wages that have been at or above $15 per hour since at least 2018 (and effectively higher than the federal minimum wage for years before). The workforce is racially diverse: industry data from the QWI show that Black workers constitute 16–22% of transportation and warehousing employment nationally, well above their 13% share of the overall labor force. This over-representation makes Amazon FC entry a particularly potent treatment for Black employment.

Labor Market Dynamics. Amazon FCs are characterized by high turnover. Annual turnover rates of 100–150% are widely reported in the trade press and corroborated by the QWI hiring flows data. This churn creates continuous demand for new workers, suggesting that FC entry represents not a one-time shock but a sustained labor demand shift. The high-turnover model also implies that the employment effects of FC entry compound over time as the facility's "gravitational pull" draws in workers from progressively wider geographic areas.

3. Data

Amazon FC Locations. I construct a panel of Amazon fulfillment center openings from the MWPVL International logistics facility database (MWPVL International, 2016), which tracks facility locations, opening dates, square footage, and facility types for major logistics operators. I retain facilities with at least 100,000 square feet or that are classified as fulfillment centers, sortation centers, or redistribution centers, excluding small delivery stations and Prime Now hubs. After geocoding addresses to county FIPS codes using the Census Bureau's ZCTA-to-county crosswalk, the treatment sample comprises 93 facilities across 63 unique counties, with first openings spanning 1997 to 2017. This is a subset of Amazon's full network, which exceeded 300 facilities by 2022; the MWPVL compilation used here covers the pre-2017 expansion reliably, after which the data moved to image-based formats. The pre-2017 period captures the initial build-out and avoids contamination from the COVID-era warehouse boom.

Quarterly Workforce Indicators. Outcome data come from the Census Bureau's Quarterly Workforce Indicators (QWI), a public-use dataset derived from the Longitudinal Employer-Household Dynamics (LEHD) program. The QWI provide county-by-quarter employment counts, average earnings, and hiring/separation flows for detailed industry-by-demographic cells. I extract NAICS 48–49 (Transportation and Warehousing) data disaggregated by race (All, White, Black, Asian) for all US counties from 2001 to 2023. The resulting panel contains 3.4 million county-quarter-race observations across 3,194 counties. As a placebo outcome, I also extract NAICS 54 (Professional, Scientific, and Technical Services) employment by race.

Treatment Definition. For each county, I define treatment as the year of the first Amazon FC opening. The 63 treated counties span 14 cohorts from 1997 to 2017, with the largest cohort in 2014 (16 counties). The remaining 3,131 counties that never received an FC serve as the control group. I aggregate the quarterly QWI data to annual frequency by averaging across quarters, as required by the Callaway-Sant'Anna estimator.

Table 1 presents pre-treatment summary statistics (2001–2007). FC counties are substantially larger than control counties: mean warehousing employment of 10,534 versus 1,570 for all races. This size differential reflects Amazon's preference for locating in counties near major metropolitan areas. Black warehousing employment averages 1,803 in FC counties versus 244 in controls, and average monthly earnings are comparable across groups ($2,536 vs. $2,488 for Black workers). The size difference is absorbed by county fixed effects in the estimation.

Table 1: Pre-Treatment Summary Statistics: Warehousing Employment by Race

	FC Counties		Control Counties	
	Mean	SD	Mean	SD
Panel A: Employment (avg. quarterly workers)				
All	10,534.2	15,185.0	1,569.7	10,462.8
White	8,107.8	11,324.3	1,235.4	8,031.9
Black	1,802.8	3,168.2	244.1	1,816.5
Asian	353.1	698.3	48.6	681.4
Panel B: Average Monthly Earnings ($)				
All	3,072	576	2,600	847
White	3,210	657	2,655	894
Black	2,536	393	2,488	1,047
Counties	60		3,113	

Notes: Pre-treatment period: 2001–2007. Employment is average quarterly beginning-of-quarter employment from the Quarterly Workforce Indicators (QWI), NAICS 48–49 (Transportation and Warehousing). Earnings are average monthly earnings. FC counties are those receiving their first Amazon fulfillment center between 1997 and 2016.

4. Empirical Strategy

I estimate the effect of Amazon FC entry on warehousing employment using the Callaway and Sant'Anna (2021) staggered difference-in-differences estimator. The identifying assumption is that, absent FC entry, warehousing employment in treated counties would have evolved in parallel with never-treated counties.

Estimation. The Callaway-Sant'Anna estimator computes group-time average treatment effects $\text{ATT}(g, t)$ for each cohort g (defined by first treatment year) and calendar year t, using never-treated counties as the comparison group. The doubly robust estimator combines outcome regression and inverse probability weighting, providing consistent estimates if either the outcome or propensity score model is correctly specified. I aggregate group-time effects into an overall ATT and dynamic event-study estimates spanning 8 years before and after treatment. Standard errors are clustered at the county level throughout.

For each race group $r \in \{\text{All, White, Black, Asian}\}$, I estimate the specification separately

on the corresponding QWI cell:

$$Y_{c,t}^{r} = \alpha_c^{r} + \gamma_t + \mathrm{ATT}^{r}(g,t) + \varepsilon_{c,t}^{r} \tag{1}$$

where $Y_{c,t}^{r}$ is log quarterly employment in county c, year t, for race r; α_c^{r} and γ_t are county-by-race and year fixed effects; and $\mathrm{ATT}^{r}(g,t)$ is the group-time treatment effect. The *racial dividend* is defined as $\mathrm{ATT}^{\mathrm{Black}} - \mathrm{ATT}^{\mathrm{White}}$. Because the Black and White ATTs are estimated from the same counties, their errors are correlated; computing the standard error of the difference as $\sqrt{\mathrm{SE}_B^2 + \mathrm{SE}_W^2}$ ignores the covariance and may overstate or understate precision. I report this conservative bound and note that a stacked specification with race interactions would provide correct joint inference; I leave this for future work.

Threats to Validity. The main threat is differential pre-trends: if counties that receive FCs were already on faster warehousing employment growth paths, the parallel trends assumption fails. I address this with (i) event-study estimates showing 8 pre-treatment coefficients close to zero, (ii) a placebo test on professional services (NAICS 54), which should be unaffected by FC entry, (iii) leave-one-cohort-out sensitivity analysis, and (iv) estimation using not-yet-treated counties as an alternative control group.

A second concern is that Amazon's FC location decisions may be endogenous to local labor market conditions. As discussed in Section 2, Houde et al. (2023) show that location timing was primarily driven by sales tax nexus laws, not local employment trends. Additionally, the event-study evidence directly tests for pre-treatment divergence.

5. Results

5.1 Main Results

Table 2 reports the main employment results. Panel A shows the Callaway-Sant'Anna ATT estimates. Overall warehousing employment increases by 0.35 log points (SE = 0.045) following FC entry. Decomposing by race reveals large heterogeneity: Black employment rises by 0.49 log points (SE = 0.064), White employment by 0.28 log points (SE = 0.037), and Asian employment by 0.64 log points (SE = 0.083). The racial dividend—the Black-White ATT differential of 0.21 log points—is statistically significant at the 1% level.

Panel B reports standard TWFE estimates for comparison. The TWFE coefficients are uniformly larger than the Callaway-Sant'Anna estimates (0.43, 0.60, and 0.33 for All, Black, and White, respectively), consistent with the well-known upward bias of TWFE in staggered designs when treatment effects grow over time (Goodman-Bacon, 2021; de Chaisemartin and

Table 2: Effect of Amazon FC Entry on Warehousing Employment by Race

	All	White	Black	Asian
Panel A: Callaway-Sant'Anna (2021)				
ATT	0.3479***	0.2816***	0.4943***	0.6405***
	(0.0452)	(0.0370)	(0.0643)	(0.0833)
Panel B: TWFE				
Post × Treated	0.4260***	0.3311***	0.5990***	—
	(0.0608)	(0.0560)	(0.0820)	
Panel C: Racial Dividend				
Black − White ATT		0.2127***		
		(0.0742)		

Notes: Dependent variable is log quarterly employment in NAICS 48–49 (Transportation and Warehousing). Panel A reports Callaway and Sant'Anna (2021) staggered DiD estimates using never-treated counties as controls. Panel B reports standard TWFE with county and year fixed effects. Panel C reports the difference between the Black and White ATTs as a test of the racial dividend hypothesis. Standard errors clustered at the county level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

D'Haultfœuille, 2020).

Panel C formalizes the racial dividend test. The Black-minus-White ATT difference of 0.21 log points (SE = 0.074) is significant at the 1% level, though the standard error assumes independence across race groups and should be interpreted as approximate (see Section 4). To benchmark magnitudes: with pre-treatment mean Black warehousing employment of 1,803 in treated counties, a 0.49 log-point effect implies roughly 1,100 additional Black workers per county. A typical Amazon FC employs 1,000–2,500 total workers, of whom approximately 20% are Black based on sector composition, suggesting 200–500 Black workers directly. The implied county-level effect exceeds direct FC employment, consistent with the treatment capturing both Amazon hiring and induced expansion of the broader logistics cluster—including third-party logistics providers, delivery contractors, and supply chain firms that co-locate near Amazon facilities.

The event-study estimates (from the dynamic aggregation) confirm clean identification. All eight pre-treatment coefficients for the overall employment specification lie within 0.03 log points of zero, with the 95% simultaneous confidence bands comfortably covering zero. The treatment effect appears immediately upon FC entry (0.06 at $t = 0$), rises to 0.22 by $t + 2$, and continues growing to 0.48 by $t + 8$. The monotonic post-treatment trajectory is consistent with Amazon's ramping-up pattern—new FCs typically hire in waves over 2–3

Table 3: Effect of Amazon FC Entry on Earnings and Hires by Race

	All	White	Black
Panel A: Log Average Monthly Earnings			
ATT	-0.0358**	-0.0286**	-0.0365**
	(0.0158)	(0.0139)	(0.0165)
Black − White		-0.0079	
		(0.0216)	
Panel B: Log All Hires			
ATT	0.6069***	0.5117***	0.4715***
	(0.0593)	(0.0499)	(0.0706)

Notes: Callaway and Sant'Anna (2021) staggered DiD estimates. Panel A: dependent variable is log average monthly earnings in NAICS 48–49. Panel B: dependent variable is log all hires. Never-treated counties as controls. Standard errors clustered at the county level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

years—and the sustained high-turnover demand creating a persistent labor market pull.

5.2 Earnings and Hiring Dynamics

Table 3 reports the effects on earnings and hiring. Panel A shows that average monthly earnings *decline* modestly following FC entry: −0.036 log points for all workers (SE = 0.016), −0.037 for Black workers (SE = 0.017), and −0.029 for White workers (SE = 0.014). The Black-White earnings differential is small and statistically indistinguishable from zero (−0.008, SE = 0.022), indicating that the earnings compression affects both groups similarly.

This earnings decline is consistent with a composition effect. Amazon's entry-level wages—around $15 per hour, or roughly $2,600 per month—are above the bottom of the warehousing wage distribution but below the sector average. As Amazon hires large numbers of entry-level workers, the sector average earnings decline even if no individual worker experiences a pay cut. This interpretation is supported by the massive increase in hiring documented in Panel B: all hires rise by 0.61 log points (SE = 0.059), reflecting Amazon's high-turnover operational model. Black hiring increases by 0.47 log points and White hiring by 0.51 log points.

The combination of large employment gains, modest earnings declines, and massive hiring increases paints a coherent picture: Amazon FC entry expands the extensive margin of employment substantially, with Black workers capturing a disproportionate share of the new positions. The intensive margin (earnings per worker) is flat to slightly negative, as Amazon's wage floor pulls down sector averages while lifting the floor for the lowest-paid workers.

Table 4: Robustness Checks

Specification	ATT	SE
Panel A: Leave-One-Cohort-Out (All Races)		
Drop 1997 cohort ($n = 1$)	0.3479	(0.0426)
Drop 1999 cohort ($n = 1$)	0.3479	(0.0430)
Drop 2000 cohort ($n = 1$)	0.3479	(0.0443)
Drop 2005 cohort ($n = 2$)	0.3591	(0.0458)
Drop 2007 cohort ($n = 1$)	0.3569	(0.0438)
Drop 2008 cohort ($n = 5$)	0.3496	(0.0448)
Drop 2010 cohort ($n = 4$)	0.3645	(0.0431)
Drop 2011 cohort ($n = 5$)	0.3526	(0.0477)
Drop 2012 cohort ($n = 8$)	0.3293	(0.0458)
Drop 2013 cohort ($n = 4$)	0.3494	(0.0455)
Drop 2014 cohort ($n = 15$)	0.3522	(0.0508)
Drop 2015 cohort ($n = 7$)	0.3293	(0.0448)
Drop 2016 cohort ($n = 3$)	0.3412	(0.0421)
Drop 2017 cohort ($n = 1$)	0.3422	(0.0449)

Notes: Panel A shows Callaway-Sant'Anna ATT estimates when each treatment cohort is dropped in turn. Dependent variable is log quarterly employment in NAICS 48–49 (all races). Never-treated counties as controls.

5.3 Robustness

Table 4 reports leave-one-cohort-out estimates for the overall employment ATT. Dropping any single treatment cohort yields an ATT between 0.33 and 0.36—a range of less than 10% of the baseline estimate. The 2012 and 2015 cohorts, when dropped, produce the smallest ATTs (0.33), while the 2005 and 2010 cohorts produce the largest (0.36). No single cohort drives the result.

Alternative Controls. Using not-yet-treated counties as the control group produces an ATT of 0.348 (SE = 0.043) for all races and 0.494 (SE = 0.060) for Black workers, virtually identical to the never-treated baseline. This rules out concerns that the never-treated group is systematically different from counties that eventually receive FCs.

Placebo Test. I estimate the same specification on NAICS 54 (Professional, Scientific, and Technical Services) employment, which should not be directly affected by FC entry. The Callaway-Sant'Anna ATT is 0.07 log points (SE = 0.037) for all races—statistically weak and economically modest compared to the 0.35 warehousing effect. The nonzero placebo is

consistent with local multiplier effects: FC workers' spending stimulates demand for local services. Importantly, the warehousing effect (0.35) is five times larger than the services spillover (0.07), confirming that the estimated employment gains are primarily sector-specific rather than reflecting general county growth.

Early vs. Late Cohorts. Splitting the sample reveals that post-2010 cohorts exhibit larger employment effects (ATT = 0.40, SE = 0.051) than pre-2010 cohorts (ATT = 0.23, SE = 0.074). This temporal heterogeneity likely reflects Amazon's increasing facility scale: post-2010 FCs are typically 800,000–1,200,000 square feet with 1,000–2,500 employees, compared to the smaller early facilities. The pattern also suggests that the racial dividend may be even larger in the most recent expansion phase, as facility size and local labor demand have grown.

6. Discussion

The racial dividend documented in this paper—a 21-log-point Black-White employment differential from Amazon FC entry—has implications that extend beyond the warehouse sector. It demonstrates that demand-side shocks to low-barrier-to-entry sectors can narrow racial employment gaps when the pre-existing racial composition of the sector favors the disadvantaged group. Black workers' over-representation in transportation and warehousing (16–22% versus 13% population share) functions as a structural amplifier: when the sector expands, Black workers are disproportionately positioned to capture the new jobs.

This mechanism—which I call the *compositional amplifier*—is distinct from both the demand-pull hypothesis (all workers benefit equally from sectoral growth) and the displacement hypothesis (new entrants crowd out incumbents). Under the compositional amplifier, differential employment incidence is a structural consequence of pre-existing occupational sorting. It does not require Amazon to engage in preferential hiring; it emerges from the intersection of sectoral demand expansion and racial occupational composition. A stronger test would interact treatment with the pre-existing county-level Black workforce share in warehousing, verifying that counties with higher ex ante Black presence see larger racial dividends; I leave this triple-difference design for future work.

The modest earnings decline (0.04 log points) raises a natural welfare question. If Amazon compresses sector earnings while expanding employment, are workers better off? The answer depends on the counterfactual. For workers at the bottom of the warehousing wage distribution, Amazon's $15 floor likely represents a wage increase. For mid-career workers at incumbents displaced by Amazon's competitive pressure, the effect may be negative. A full welfare analysis requires linked employer-employee data that the QWI's cell structure cannot

provide; this is a natural extension for future work using the LEHD microdata.

Two limitations deserve emphasis. First, the treatment sample covers only 63 counties through 2016. Amazon's post-2016 expansion—adding roughly 200 facilities during the COVID-era warehouse boom—likely generated additional racial employment effects that this paper cannot capture. Extending the facility database through 2022 would substantially increase power and allow testing whether the racial dividend persists at scale. Second, the QWI's county-level aggregation cannot identify within-county displacement: if Amazon pulls workers from other logistics employers, the net employment effect overstates welfare gains. The event-study's monotonic post-treatment growth suggests that expansion dominates displacement, but firm-level analysis would sharpen this conclusion.

7. Conclusion

Amazon's fulfillment center expansion generated differential racial employment incidence: Black warehousing employment grew substantially faster than White employment in counties receiving FCs, consistent with the structural intersection of sectoral demand growth and pre-existing racial occupational composition. Whether this compositional amplifier constitutes a welfare gain for Black workers depends on counterfactual employment and the earnings distribution—questions that county-level data cannot fully resolve. What the evidence does establish is that demand-side shocks to low-barrier-to-entry sectors have unequal racial incidence, a pattern that should inform debates about industrial policy, warehouse zoning, and the distributional consequences of Amazon's labor market footprint.

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

Autor, David, David Dorn, Lawrence F. Katz, Christina Patterson, and John Van Reenen, "The Fall of the Labor Share and the Rise of Superstar Firms," *Quarterly Journal of Economics*, 2020, *135* (2), 645–709.

Autor, David H., David Dorn, and Gordon H. Hanson, "The China Syndrome: Local Labor Market Effects of Import Competition in the United States," *American Economic Review*, 2013, *103* (6), 2121–2168.

Callaway, Brantly and Pedro H. C. Sant'Anna, "Difference-in-Differences with Multiple Time Periods," *Journal of Econometrics*, 2021, *225* (2), 200–230.

Charles, Kerwin Kofi, Erik Hurst, and Mariel Schwartz, "The Transformation of Manufacturing and the Decline in US Employment," *NBER Macroeconomics Annual*, 2018, *33*, 307–372.

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

Derenoncourt, Ellora, "Can You Move to Opportunity? Evidence from the Great Migration," *American Economic Review*, 2022, *112* (2), 369–408.

Dix-Carneiro, Rafael and Brian K. Kovak, "Trade Liberalization and Regional Dynamics," *American Economic Review*, 2017, *107* (10), 2908–2946.

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

Houde, Jean-François, Peter Newberry, and Katja Seim, "Nexus Tax Laws and Economies of Density in E-Commerce: A Study of Amazon's Fulfillment Center Network," *Econometrica*, 2023, *91* (1), 147–190.

Jones, Janelle and Ben Zipperer, "Unfulfilled Promises: Amazon Fulfillment Centers Do Not Generate Broad-Based Employment Growth," Technical Report, Economic Policy Institute 2018.

MWPVL International, "Amazon Global Supply Chain and Fulfillment Center Network," 2016. Accessed via archived compilation, 2016.

A. Data Appendix

Amazon FC Locations. Facility data were obtained from the MWPVL International logistics database via a publicly archived compilation (MWPVL International, 2016). The raw dataset contains 156 North American entries. I retain US facilities, drop Prime Now hubs (codes beginning with "U"), Fresh delivery stations (codes beginning with "D"), and facilities below 100,000 square feet that are not classified as fulfillment, sortation, or redistribution centers. Opening years are extracted from date strings in the "YearOpened" field using regular expressions matching four-digit years in the 1990–2020 range. Addresses are geocoded to county FIPS codes using the Census Bureau's 2020 ZCTA-to-county crosswalk (primary) and OpenStreetMap/Nominatim geocoding (secondary), followed by spatial point-in-polygon assignment to Census county boundaries. The final sample contains 93 facilities across 63 unique counties.

QWI Data. The Quarterly Workforce Indicators are produced by the Census Bureau's LEHD program from state unemployment insurance records matched to Census demographic data. I query the race-by-industry ("rh/ns") QWI files for NAICS 48–49 (Transportation and Warehousing) and NAICS 54 (Professional Services), restricting to private-sector employment (owner code A05) and the four main racial categories: All (A0), White (A1), Black (A2), and Asian (A4). The data cover 2001–2023 across 3,194 counties. Quarterly observations are averaged to annual frequency.

Variable Definitions.

- **Employment (Emp):** Beginning-of-quarter employment count—the number of workers with positive earnings in both the current and previous quarters.
- **Average Earnings (EarnS):** Average monthly earnings for stable jobs (workers present at both the beginning and end of the quarter).
- **All Hires (HirA):** Workers with positive earnings in the current quarter who did not have positive earnings with the same employer in the previous quarter.
- **Separations (Sep):** Workers with positive earnings in the current quarter who do not have positive earnings with the same employer in the following quarter.

B. Robustness Appendix

Pre-Trend Assessment. The event-study estimates for the overall employment specification show all eight pre-treatment coefficients ($e = -8$ to $e = -1$) lie within the range [−0.006, 0.024], with none individually or jointly significant at the 5% level using simultaneous confidence bands. The largest pre-treatment coefficient (0.024 at $e = -6$) is an order of magnitude smaller than the first post-treatment coefficient (0.057 at $e = 0$).

Leave-One-Cohort-Out Detail. The 14 leave-one-cohort-out estimates reported in Table 4 show remarkable stability. The range of estimates (0.329–0.365) implies a coefficient of variation of less than 3%. The 2012 and 2015 cohorts produce the largest shifts when dropped, each reducing the ATT to 0.329—still well within the 95% confidence interval of the baseline estimate of 0.348.

C. Standardized Effect Sizes

Table 5: Standardized Effect Sizes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Panel A: Pooled						
Employment (All)	0.3479	0.0452	2.0924	0.1663	0.0216	Large positive
Earnings (All)	-0.0358	0.0158	0.3353	-0.1068	0.0470	Moderate negative
Hires (All)	0.6069	0.0593	1.9055	0.3185	0.0311	Large positive
Panel B: Heterogeneous (Race)						
Employment (Black)	0.4943	0.0643	2.3252	0.2126	0.0277	Large positive
Employment (White)	0.2816	0.0370	2.0598	0.1367	0.0179	Moderate positive

Notes: **Country:** United States. **Research question:** Does Amazon fulfillment center entry disproportionately increase Black employment and earnings in county-level warehousing labor markets? **Policy mechanism:** Amazon's staggered opening of large fulfillment centers (typically 500,000–1,200,000 sq ft) in US counties creates discrete demand shocks for logistics workers; these positions require no college degree and draw from a labor pool where Black workers are over-represented (16–22% of NAICS 48–49 employment). **Outcome definition:** Log quarterly beginning-of-quarter employment in NAICS 48–49 (Transportation and Warehousing) from the Census Bureau's Quarterly Workforce Indicators (QWI), disaggregated by race. **Treatment:** Binary: county receives its first Amazon fulfillment center (staggered, 1997–2016). **Data:** QWI race panel (county × quarter × race), NAICS 48–49, 2001–2023; Amazon FC locations from MWPVL International compilation. **Method:** Callaway and Sant'Anna (2021) staggered DiD with never-treated controls; standard errors clustered at the county level; 63 treated counties, 3133+ control counties. **Sample:** US counties with non-missing QWI warehousing employment and at least 15 years of panel data; excludes counties with missing race-specific employment throughout. SDE = $\hat{\beta}$/SD(Y) where SD(Y) is the pre-treatment 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).