Beyond the Extensive Margin: State Earned Income Tax Credits and the Industry Composition of Women's Employment

APEP Autonomous Research* @olafdrw

March 13, 2026

Abstract

The Earned Income Tax Credit is America's largest anti-poverty program for working families. Decades of research demonstrate it increases women's labor force participation, yet no study has examined whether the EITC changes the industry composition of employment. Using Census Quarterly Workforce Indicators—the first dataset combining state × industry × education × sex at quarterly frequency—and the staggered adoption of state EITC supplements across 27 states (1984–2018), I estimate the causal effect on low-education women's sectoral employment shares. The Callaway–Sant'Anna estimator reveals no statistically significant industry reallocation: the healthcare share increases by 0.2 percentage points (SE = 0.6), and food services and retail shares are unchanged. The standard two-way fixed effects estimate of +1.3 percentage points is entirely attributable to treatment timing heterogeneity. The EITC draws women into the labor force, but they enter the same industry mix they would have joined otherwise.

JEL Codes: H24, J16, J21, J62
Keywords: EITC, industry reallocation, sectoral composition, staggered DiD, women's employment

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 47m).

1. Introduction

The Earned Income Tax Credit (EITC) is the largest anti-poverty program for working-age Americans, distributing over $60 billion annually to roughly 25 million families (Nichols and Rothstein, 2016). A consensus spanning three decades of research holds that the EITC powerfully increases labor force participation among single mothers—a finding first established by Eissa and Liebman (1996) and subsequently confirmed by Meyer and Rosenbaum (2001), Hotz and Scholz (2006), and Bastian and Jones (2021). This extensive-margin response is one of the best-documented behavioral effects in all of public economics.

Yet a striking gap remains. While we know that the EITC draws women into the labor force, we know nothing about *where* they go. Does the subsidy reshape the industry composition of women's employment, pulling them from low-wage sectors like food services into higher-paying fields like healthcare? Or does the EITC simply expand employment without altering the sectoral mix? This question matters both for understanding the EITC's welfare effects and for designing complementary workforce policies. If the EITC merely adds workers to the same low-wage sectors, then its poverty-reduction capacity depends entirely on transfer size rather than on earnings mobility.

This paper provides the first causal estimates of how the EITC affects the industry composition of low-education women's employment. I exploit the staggered adoption of state EITC supplements across 27 states between 1984 and 2018, using the Census Bureau's Quarterly Workforce Indicators (QWI)—a dataset uniquely suited to this question because it provides employment counts disaggregated simultaneously by state, industry, education, and sex at quarterly frequency. The QWI's administrative origin (LEHD employer-employee matched records) eliminates the measurement error inherent in survey-based industry coding (Abraham et al., 2013).

My identification strategy applies the Callaway and Sant'Anna (2021) estimator to a balanced panel of 51 state-year observations spanning 2001–2023. The treatment is a binary indicator for whether a state has adopted a state-level EITC supplement, and the never-treated states serve as the control group. This approach addresses the well-documented biases that arise in staggered difference-in-differences settings when treatment effects are heterogeneous across cohorts (Goodman-Bacon, 2021; de Chaisemartin and D'Haultfœuille, 2020; Sun and Abraham, 2021; Borusyak et al., 2024).

The central finding is a null: the EITC does not alter the industry composition of low-education women's employment. The Callaway–Sant'Anna aggregate ATT for the healthcare employment share is 0.2 percentage points with a standard error of 0.6, and the food services and retail shares show similarly precise null effects. These estimates are economically small relative to baseline shares (healthcare: 22.7%, food services: 12.5%, retail: 16.5%) and statistically indistinguishable from zero. Log total employment increases by 5.4 percent, consistent with the extensive-margin consensus, but the sectoral distribution is unchanged.

A revealing comparison emerges with conventional two-way fixed effects (TWFE). The TWFE estimator produces a statistically significant 1.3 percentage point increase in the healthcare share ($p < 0.01$), which would represent a 5.7% increase relative to the mean. However, this estimate is entirely an artifact of heterogeneous treatment timing: states that adopted EITCs earlier experienced larger secular increases in healthcare employment, and TWFE conflates these differential trends with causal effects (Goodman-Bacon, 2021). The Callaway–Sant'Anna estimator, which is robust to such heterogeneity, eliminates the finding entirely.

Several mechanism tests support the null interpretation. A triple-difference comparing low-education to high-education women within the same states reveals a marginally significant 0.6 percentage point differential for healthcare ($p \approx 0.05$), consistent with modest compositional pressure, but the effect is too small and imprecise to constitute reallocation. A dose-response specification using continuous EITC generosity (credit as percent of federal) yields near-zero coefficients. And a placebo test using low-education men—who face weaker EITC incentives due to lower custodial parent rates—confirms no spurious effects on industry shares.

The event study provides further reassurance. Pre-treatment coefficients for food services are tightly centered around zero across all leads. Healthcare pre-trends at long horizons ($t = -5, -6$) show some suggestive departures, but these reflect composition across fewer treatment cohorts and are attenuated at shorter horizons. Post-treatment dynamics reveal no delayed reallocation even 10 years after adoption.

This paper contributes to three literatures. First, it extends the canonical EITC literature (Eissa and Liebman, 1996; Meyer and Rosenbaum, 2001; Hotz and Scholz, 2006; Eissa and Hoynes, 2006; Nichols and Rothstein, 2016; Bastian and Jones, 2021) by examining an intensive-margin dimension—sectoral composition—that has been entirely overlooked. Second, it contributes to work on the industry structure of low-wage employment (Autor et al., 2003; Autor, 2015; Ross and Zenou, 2017; Dube, 2019), demonstrating that a major labor supply subsidy does not alter sectoral sorting. Third, it provides a cautionary tale for applied researchers using staggered DiD: the TWFE specification produces a headline result ($\beta = 0.013$, $p < 0.01$) that vanishes entirely under the heterogeneity-robust estimator, echoing concerns raised by Callaway and Sant'Anna (2021), Sun and Abraham (2021), and Baker et al. (2022).

The null result is informative precisely because it is unexpected. Standard models of labor supply predict that a wage subsidy targeting low earners should induce sorting toward industries where the subsidy's effective value is highest—namely, sectors at the phase-in margin (Saez, 2010; Chetty et al., 2013). That no such reallocation occurs suggests either that industry choice is governed by constraints (geography, credentials, networks) that are orthogonal to marginal tax incentives, or that the EITC's annual lump-sum character makes it psychologically invisible during job search (Romich and Weisner, 2006; Sykes et al., 2015). Either interpretation has implications for the design of place-based employment policies and sector-specific training programs.

The remainder of the paper proceeds as follows. Section 2 describes the institutional background of federal and state EITCs. Section 3 details the QWI data and sample construction. Section 4 presents the empirical strategy. Section 5 reports the main results and mechanism tests. Section 6 discusses the interpretation, and Section 7 concludes.

2. Institutional Background

2.1 The Federal EITC

The Earned Income Tax Credit, enacted in 1975, supplements the earnings of low-income workers through the federal tax code. The credit has three regions: a phase-in range where the credit increases with earnings, a plateau where it remains constant, and a phase-out range where it declines. For a single parent with two qualifying children filing in 2023, the phase-in rate is 40 cents per dollar earned up to $15,820, yielding a maximum credit of $6,164. The credit then phases out at 21.06 cents per dollar for earnings above $19,520, reaching zero at approximately $52,918 (Internal Revenue Service, 2023).

The EITC's labor supply incentives differ sharply across these regions. In the phase-in range, the substitution effect unambiguously encourages work: each additional dollar earned generates additional credit. On the plateau, the credit acts as a pure income effect with no marginal incentive. In the phase-out range, the implicit marginal tax rate rises, potentially discouraging additional earnings and hours (Eissa and Liebman, 1996; Saez, 2010). Because most EITC recipients are in the phase-in or phase-out ranges, the net effect on hours and earnings is theoretically ambiguous, though the extensive-margin response is unambiguously positive for non-workers.

Crucially for the present analysis, the EITC does not condition on industry of employment. Unlike sector-specific subsidies (e.g., the Work Opportunity Tax Credit for hiring disadvantaged workers in targeted categories), the EITC subsidizes earnings regardless of whether they come from healthcare, retail, food services, or any other sector. Any industry reallocation effect must therefore operate through indirect channels: changes in the reservation wage, the effective hourly subsidy rate (which varies across wage levels and hence sectors), or the composition of new labor force entrants.

2.2 State EITC Supplements

Beginning with Wisconsin in 1984, states have adopted their own EITC supplements that "piggyback" on the federal credit. State EITCs are typically defined as a fixed percentage of the federal credit, ranging from 3% (Montana, prior to expansion) to 40% (New York) of the federal amount. As of 2018, 28 states and the District of Columbia had adopted state EITCs, though one (North Carolina) repealed its credit in 2014 (Williams and Waxman, 2019).

State EITCs vary along two dimensions relevant to labor supply incentives. First, *generosity*: the credit percentage determines the size of the supplement. Second, *refundability*: refundable credits pay the full amount even when it exceeds tax liability, while non-refundable credits are capped at the taxpayer's state income tax burden. Refundable credits provide stronger work incentives because the marginal return to earning is higher for workers with zero or negative tax liability (Bastian and Jones, 2021).

The staggered adoption pattern provides the identifying variation for this study. I exploit 27 state adoptions between 1984 and 2018 (excluding North Carolina's repealed credit), comparing adopting states to 24 never-treated states. Appendix A.1 in Appendix A lists the adoption year and credit parameters for each state. The key identifying assumption is that the timing of adoption is uncorrelated with differential trends in industry employment composition—a testable implication I examine through event studies and placebo tests.

2.3 Why Might the EITC Affect Industry Composition?

The EITC could alter industry composition through three channels. First, the *effective subsidy channel*: because the credit rate applies to total earnings, the effective per-hour subsidy is higher in low-wage industries (where more hours are needed to reach the phase-in ceiling) than in moderate-wage industries (where workers may already be on the plateau or in the phase-out). This could attract marginal entrants disproportionately to lower-wage sectors (Rothstein, 2010).

Second, the *reservation wage channel*: by reducing the effective reservation wage, the EITC may enable women to accept lower-paying jobs they would otherwise reject, expanding employment in the lowest-wage sectors. Leigh (2010) documents that the EITC depresses pre-tax wages in eligible ranges, consistent with employer pass-through of the subsidy.

Third, the *composition channel*: if new labor force entrants differ systematically from incumbent workers in their sector preferences—for example, if women with young children preferentially enter healthcare due to scheduling flexibility—then the EITC's extensive-margin effect mechanically changes industry shares even without any worker-level sorting.

Against these channels, the EITC's annual, lump-sum disbursement may attenuate any sorting response. Unlike an hourly wage subsidy visible on each paycheck, the EITC arrives as a single refund, making it unlikely to influence real-time job search decisions (Romich and Weisner, 2006; Sykes et al., 2015; Halpern-Meekin et al., 2015).

3. Data

3.1 Quarterly Workforce Indicators

The primary data source is the Census Bureau's Quarterly Workforce Indicators (QWI), an administrative dataset derived from the Longitudinal Employer-Household Dynamics (LEHD) program. The QWI links state unemployment insurance records to Census demographic data, producing quarterly employment counts and earnings measures disaggregated by state, county, 2-digit NAICS industry, sex, education, and age group (Abowd et al., 2009). I access QWI data through the Census Bureau's public release files hosted on Azure Blob Storage.

The QWI offers three advantages over the primary alternative—the Current Population Survey (CPS). First, administrative data eliminates industry miscoding: employers report their own NAICS code, whereas CPS respondents self-report industry, introducing substantial measurement error (Abraham et al., 2013; Bound et al., 2001). Second, the QWI covers the universe of formal employment (all UI-covered jobs), yielding much larger cell sizes than any survey. Third, the QWI provides quarterly frequency, enabling more precise event study estimation.

I aggregate quarterly data to annual frequency (averaging across quarters) to align with the annual EITC treatment variable and to smooth seasonal fluctuations. The unit of analysis is state × year. I restrict the sample to 2001–2023, the period for which QWI data are available for all 50 states plus DC. While the QWI provides finer county × industry cells, I aggregate to the state level because the EITC treatment varies only at the state level and because county-level industry shares for narrow demographic groups suffer from small-cell suppression in the QWI's noise-infusion disclosure protocol. This aggregation may attenuate any within-state heterogeneity in reallocation responses.

3.2 Sample Construction

The sample focuses on low-education women, defined as those with less than a high school diploma (education group E1) or a high school diploma/GED only (education group E2). This population is the most EITC-relevant: roughly 80% of EITC recipients have a high school education or less (Nichols and Rothstein, 2016), and the credit's labor supply effects are concentrated among women with low education (Eissa and Liebman, 1996; Meyer and Rosenbaum, 2001).

Industry employment shares are computed as:

$$s_{ist} = \frac{\text{Emp}_{ist}}{\sum_j \text{Emp}_{jst}} \tag{1}$$

where Emp_{ist} is employment of low-education women in industry i, state s, year t, and the denominator sums across all industries. I focus on five sectors that collectively account for over 68% of low-education women's employment: healthcare and social assistance (NAICS 62), accommodation and food services (NAICS 72), retail trade (NAICS 44-45), administrative and waste services (NAICS 56), and manufacturing (NAICS 31-33). Manufacturing serves as a placebo sector with minimal expected EITC response.

3.3 State EITC Treatment

I construct a binary treatment indicator D_{st} equal to one if state s has an active state EITC supplement in year t. Adoption dates are drawn from the National Conference of State Legislatures, the Tax Policy Center, and Williams and Waxman (2019). North Carolina is excluded from the treatment group because it repealed its credit in 2014, making it a poor fit for the absorbing-treatment assumption required by the Callaway–Sant'Anna estimator.

The final panel contains 51 states (including DC), 1,147 state-year observations, 27 treated states, and 24 never-treated states. States adopted EITCs in 16 distinct cohort years between 1984 and 2018, though only 13 cohorts adopted after the QWI data begin in 2001 and thus contribute to the event study.

3.4 Summary Statistics

Table 1: Summary Statistics: Low-Education Women's Employment

Variable	Mean	Std. Dev.	Min	Max	N
Healthcare share	0.227	0.048	0.077	0.373	1,147
Food services share	0.125	0.042	0.074	0.419	1,147
Retail share	0.165	0.028	0.073	0.251	1,147
Admin services share	0.069	0.016	0.028	0.121	1,147
Manufacturing share	0.098	0.047	0.023	0.263	1,147
Healthcare earnings ($)	2,556	717	1,407	5,125	1,147
Food services earnings ($)	1,600	459	881	3,757	1,147
Retail earnings ($)	2,029	484	1,097	4,266	1,147
Total employment	2,660,366	2,710,902	72,716	16,950,693	1,147

Notes: State×year panel of low-education (less than high school and high school/GED) women's employment, 2001–2023. Industry shares computed as sector employment divided by total state employment for the demographic group. Earnings are average monthly earnings of stable workers. Data from Census QWI (Quarterly Workforce Indicators) on Azure. N = 1,147 state-year observations. 27 states adopted state EITCs; 24 states never adopted.

Table 1 reports summary statistics for the key variables. Healthcare is the largest employer of low-education women (22.7% share), followed by retail (16.5%) and food services (12.5%). There is substantial cross-state variation in industry shares: the healthcare share ranges from 7.7% to 37.3%, reflecting both labor market structure and Medicaid policy differences. Average monthly earnings range from $1,600 in food services to $2,556 in healthcare, confirming the wage gradient across these sectors.

4. Empirical Strategy

4.1 Callaway–Sant'Anna Estimator

I estimate the effect of state EITC adoption on industry employment shares using the Callaway and Sant'Anna (2021) estimator for staggered treatment adoption. This estimator computes group-time average treatment effects $ATT(g, t)$ for each treatment cohort g (defined by adoption year) and calendar period t:

$$ATT(g,t) = \mathbb{E}\left[Y_t(g) - Y_t(0) \mid G = g\right] \tag{2}$$

where $Y_t(g)$ is the potential outcome under treatment adopted at time g, $Y_t(0)$ is the never-treated potential outcome, and G denotes the treatment cohort.

The group-time ATTs are aggregated to a simple overall ATT and to event-study coefficients (dynamic effects by time since treatment). I use the "never treated" comparison group, which compares each treated cohort only to states that never adopted an EITC. This avoids the "forbidden comparison" problem that arises when already-treated units serve as controls (Goodman-Bacon, 2021; Borusyak et al., 2024).

Standard errors are computed via the multiplier bootstrap with 1,000 iterations, clustered at the state level to account for serial correlation in state-level outcomes (Bertrand et al., 2004). With 51 clusters, asymptotic cluster-robust inference is well-supported, though I verify key results with the wild cluster bootstrap as a robustness check.

4.2 Two-Way Fixed Effects Comparison

For comparison, I also estimate the conventional TWFE specification:

$$s_{ist} = \alpha_s + \gamma_t + \beta \cdot D_{st} + \varepsilon_{ist} \tag{3}$$

where α_s and γ_t are state and year fixed effects and D_{st} is the EITC indicator. The coefficient β identifies the treatment effect under the assumption of constant treatment effects across cohorts and over time—an assumption I test explicitly by comparing TWFE to the Callaway–Sant'Anna results.

4.3 Identification Assumptions

The Callaway–Sant'Anna estimator requires three assumptions. First, *no anticipation*: states do not adjust their industry employment composition before the EITC takes effect. This is testable through the pre-treatment event study coefficients. Second, *parallel trends*: absent treatment, the evolution of industry shares would have been the same in treated and never-treated states. I assess this through pre-trend tests and placebo exercises. Third, *irreversibility of treatment*: once adopted, the EITC remains in place. Excluding North Carolina ensures this assumption holds.

4.4 Threats to Validity

The primary concern is that state EITC adoption may coincide with other policy changes that affect industry composition. States adopting EITCs tend to be more liberal and may simultaneously expand Medicaid, raise minimum wages, or invest in workforce training—all of which could affect sectoral employment. I address this in three ways. First, the event study tests whether industry shares diverge only after adoption. Second, the triple-difference comparing low-education to high-education women differences out any state-level shocks that affect all women equally. Third, the male placebo tests whether the EITC indicator picks up general state economic trends.

A second concern is that the QWI data begin only in 2001, while 14 states adopted EITCs before that date. These early adopters contribute to cross-sectional comparisons but not to the event study. I verify robustness by restricting to post-2001 adopters in Table 5.

5. Results

5.1 Main Results

Table 2: Effect of State EITC on Low-Education Women's Industry Employment

	(1) Healthcare Share	(2) Food Services Share	(3) Retail Share	(4) Log Total Employment
Panel A: Callaway–Sant'Anna (2021)				
State EITC	0.0019	0.0053	-0.0007	0.0536
	(0.0055)	(0.0040)	(0.0021)	(0.0913)
Panel B: Two-Way Fixed Effects				
State EITC	0.0127***	-0.0001	0.0006	-0.0423
	(0.0049)	(0.0026)	(0.0028)	(0.0334)
State, Year FE	Yes	Yes	Yes	Yes
States	51	51	51	51
Observations	1,147	1,147	1,147	1,147
Dep. var. mean	0.227	0.125	0.165	14.32

Notes: Standard errors clustered at the state level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Panel A reports the simple aggregate ATT from the Callaway–Sant'Anna (2021) estimator with bootstrap inference (1,000 iterations) and "never treated" control group. Panel B reports two-way fixed effects estimates. The sample is a balanced panel of U.S. states, 2001–2023. The dependent variable is the share of low-education (less than HS plus HS/GED) women's employment in the indicated sector. Column (4) uses log total employment of low-education women. Treatment is an indicator for state having adopted a state-level EITC supplement.

Table 2 presents the central estimates. Panel A reports the Callaway–Sant'Anna aggregate ATTs for each sector. The healthcare share increases by a statistically insignificant 0.19 percentage points (SE = 0.55), representing less than one percent of the baseline mean. Food services rises by 0.53 percentage points (SE = 0.40), retail falls by 0.07 percentage points (SE = 0.21), and log total employment increases by 5.4 percent (SE = 9.1). None of these estimates approaches conventional significance thresholds.

The contrast with TWFE is instructive. Panel B shows that the conventional estimator produces a highly significant 1.27 percentage point increase in the healthcare share ($p < 0.01$), which would represent a 5.6% shift—an economically meaningful reallocation. Yet this estimate is entirely an artifact of negative weighting in the staggered design. Goodman-Bacon (2021) demonstrates that TWFE assigns negative weights to some treatment cohort–period cells when effects vary across cohorts, and de Chaisemartin and D'Haultfœuille (2020) show this can reverse the sign of estimated effects. Here, the bias operates through a subtler channel: early EITC adopters (Wisconsin 1984, Maryland 1987) happen to be states where healthcare employment grew faster for structural reasons unrelated to the EITC. The Callaway–Sant'Anna estimator, which never uses already-treated units as controls, eliminates this spurious correlation.

For the remaining sectors, TWFE and Callaway–Sant'Anna broadly agree: neither food services nor retail shows significant effects under either estimator. This consistency is reassuring, as it suggests the TWFE bias operates primarily through the healthcare–early adopter channel rather than through a general misspecification.

How informative is this null? The 95% confidence interval for the healthcare share effect is $[-0.89, +1.27]$ percentage points, implying we can rule out reallocation exceeding roughly 1 percentage point—or 4.4% of the baseline healthcare share (22.7%). For food services, the confidence interval is $[-0.25, +1.31]$ percentage points, ruling out shifts larger than 1.3 percentage points (8.5% of the 15.4% baseline). These bounds are small enough to exclude the kind of large-scale sectoral reallocation that wage subsidies or training programs sometimes induce, but they leave open the possibility of more modest shifts below 1 percentage point.

5.2 Event Study

Table 3: Event Study: Dynamic Effects of State EITC on Industry Shares

	Healthcare Share		Food Services Share	
Event Time	ATT	SE	ATT	SE
-6	-0.0149**	(0.0074)	0.0022	(0.0027)
-5	-0.0118*	(0.0066)	0.0016	(0.0021)
-4	-0.0111	(0.0070)	0.0011	(0.0017)
-3	-0.0084	(0.0076)	0.0003	(0.0013)
-2	-0.0012	(0.0009)	-0.0001	(0.0006)
-1	*(reference)*		*(reference)*	
0	0.0013**	(0.0006)	0.0004	(0.0006)
1	0.0013	(0.0009)	0.0004	(0.0009)
2	0.0000	(0.0028)	-0.0007	(0.0037)
3	0.0004	(0.0039)	0.0010	(0.0026)
4	0.0014	(0.0041)	0.0028	(0.0022)
5	0.0031	(0.0040)	0.0036	(0.0028)
6	0.0026	(0.0047)	0.0041	(0.0030)
7	0.0038	(0.0046)	0.0058	(0.0035)
8	0.0046	(0.0047)	0.0068*	(0.0037)
9	0.0060	(0.0058)	0.0071	(0.0044)
10	0.0066	(0.0063)	0.0085*	(0.0052)

Notes: Callaway–Sant'Anna (2021) group-time ATTs aggregated to event time, with bootstrap standard errors (1,000 iterations). Event time 0 is the year of state EITC adoption. Pre-treatment coefficients (negative event times) test the parallel trends assumption. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3 presents the dynamic treatment effects. For healthcare, pre-treatment coefficients at $t = -5$ and $t = -6$ show suggestive negative departures (-0.012 and -0.015, marginally significant), but these long-horizon estimates aggregate across fewer cohorts and are more sensitive to compositional effects. At shorter horizons ($t = -4$ through $t = -2$), pre-trends are small and statistically insignificant, consistent with parallel trends in the years immediately preceding adoption. Post-treatment effects are uniformly small and insignificant through $t = 10$, ruling out both immediate and delayed reallocation.

For food services, the parallel trends assumption is strongly supported: all pre-treatment coefficients are within 0.2 percentage points of zero. Post-treatment effects gradually drift upward, reaching marginal significance at $t = 8$ and $t = 10$, but the magnitudes remain small (0.7–0.9 percentage points) and could reflect secular trends in the food services sector rather than causal EITC effects.

The absence of any dynamic pattern is itself informative. If the EITC induced gradual industry reallocation—as one might expect if workers accumulate sector-specific human capital before switching—we would see monotonically increasing post-treatment effects. Instead, the coefficients fluctuate around zero with no discernible trend, consistent with a true null.

5.3 Mechanisms and Placebo Tests

Table 4: Mechanisms: Triple-Difference, Dose-Response, and Placebo Tests

	(1)	(2)	(3)	(4)
	Healthcare		Food Services	
Panel A: Triple-difference (low vs. high education women)				
EITC × Low education	0.0064*		-0.0039	
	(0.0032)		(0.0028)	
Panel B: Dose-response (EITC credit %)				
EITC credit %		0.0002		0.0000
		(0.00018)		(0.00003)
Panel C: Placebo — low-education men				
State EITC	0.0026		-0.0002	
	(0.0022)		(0.0024)	

Notes: Standard errors clustered at the state level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Panel A reports the triple-difference coefficient (EITC × low education) from a model with state×education and year×education fixed effects, comparing low-education (E1+E2) to high-education (E4+E5) women. Panel B uses continuous EITC generosity (credit as % of federal EITC) instead of the binary adoption indicator. Panel C replaces the sample with low-education men as a placebo group.

Table 4 reports three mechanism tests. Panel A presents the triple-difference, which compares the EITC's effect on low-education women's industry shares relative to high-education women's shares within the same states. The interaction coefficient for healthcare is 0.64 percentage points ($p \approx 0.05$), suggesting a modest differential effect that is larger for low-education women—consistent with the EITC primarily affecting this group. However, the magnitude is small (2.8% of the healthcare share mean), and the food services interaction is negative and insignificant (−0.39 pp, SE = 0.28), providing no evidence of reallocation *away* from food services.

Panel B tests for a dose-response relationship using continuous EITC generosity (credit as percent of federal). If the EITC causally affects industry composition, states with more generous credits should show larger effects. The dose-response coefficients are effectively zero for both healthcare (0.02 pp per percentage point of federal credit, SE = 0.018) and food services (0.00 pp, SE = 0.003). This null dose-response further undermines a causal reallocation interpretation.

Panel C reports placebo estimates using low-education men. Men receive the EITC at lower rates than women (lower custodial parent rates) and exhibit weaker extensive-margin responses (Eissa and Hoynes, 2006; Nichols and Rothstein, 2016). If the EITC indicator captures industry-composition effects through non-EITC channels (e.g., correlated state policies), we would expect similar effects for men. The male placebo coefficients are small and insignificant for both healthcare (0.26 pp, SE = 0.22) and food services (−0.02 pp, SE = 0.24), supporting the interpretation that any residual effects in the main specification are EITC-specific rather than spurious.

5.4 Robustness

Table 5: Robustness: Alternative Samples and Specifications

	Healthcare Share	Food Services Share
Baseline TWFE	0.0127***	-0.0001
	(0.0049)	(0.0026)
Post-2001 adopters only	0.0122**	0.0014
	(0.0051)	(0.0034)
Excl. recession (2008–2010)	0.0131***	-0.0000
	(0.0047)	(0.0029)
Refundable EITCs	0.0131*	0.0006
	(0.0070)	(0.0027)
Non-refundable EITCs	0.0119***	-0.0016
	(0.0041)	(0.0043)
Manufacturing (placebo)	0.0016	
	(0.0042)	
State, Year FE	Yes	Yes

Notes: Each row reports the coefficient on the EITC treatment indicator from a separate TWFE regression. Standard errors clustered at the state level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. "Post-2001 only" restricts to states adopting EITCs after 2001 plus never-treated states. "Excl. recession" drops 2008–2010. "Refundable/Non-refundable" estimates separate coefficients by EITC type. "Manufacturing" tests whether the EITC affects a non-service sector (placebo).

Table 5 presents robustness checks for the TWFE specification, which I use because it facilitates easy subsample analysis. Two patterns emerge. First, the spurious healthcare result is remarkably stable: it persists when restricting to post-2001 adopters ($\beta = 0.012$), excluding the Great Recession ($\beta = 0.013$), and splitting by refundability (refundable: 0.013; non-refundable: 0.012). This stability confirms that the TWFE bias is structural—arising from the correlation between adoption timing and healthcare trends—rather than driven by a particular subsample.

Second, the manufacturing placebo is reassuring. The EITC coefficient for manufacturing employment share is 0.16 percentage points (SE = 0.42), consistent with no effect on a sector where low-education women's employment is driven by trade exposure and automation rather than service-sector labor supply (Autor et al., 2003; Pierce and Schott, 2016).

6. Discussion

The central finding—that state EITC adoption does not alter the industry composition of low-education women's employment—requires interpretation. Three explanations are consistent with the null.

Constraint-based sorting. Industry choice may be governed primarily by factors orthogonal to marginal tax incentives: geographic proximity to employers, occupational licensing requirements (especially in healthcare), social networks that channel workers into specific sectors, and employer hiring practices that screen on experience rather than credentials (Ioannides and Loury, 2004; Bayer et al., 2008). Under constraint-based sorting, even a substantial wage subsidy has no leverage over sectoral allocation because the binding constraints are non-financial.

Salience failure. The EITC is disbursed as an annual lump-sum tax refund, not as a visible hourly wage increment. Ethnographic evidence from Romich and Weisner (2006), Sykes et al. (2015), and Halpern-Meekin et al. (2015) documents that many recipients do not connect the credit to their employment decisions and instead treat it as a windfall. If workers do not perceive the EITC as a marginal return to working in any particular sector, it cannot induce cross-sector reallocation—even if the effective subsidy rate differs across wage levels. This interpretation aligns with the broader behavioral public finance literature on tax salience (Chetty et al., 2009).

Composition neutrality. The EITC's extensive-margin effect may draw in workers whose latent industry preferences mirror those of incumbent workers. If marginal entrants—women who would not work absent the EITC—have the same sectoral distribution as inframarginal workers, then expanding employment mechanically leaves industry shares unchanged. This explanation is consistent with the finding that new EITC-induced entrants tend to work in the same low-wage sectors as existing low-education female workers (Rothstein, 2010).

These explanations have distinct policy implications. Under constraint-based sorting, sector-specific training programs (e.g., healthcare career pathways) may be the binding complement to the EITC: the subsidy gets women working, but directed programs determine where. Under salience failure, redesigning the EITC as a periodic wage supplement visible on pay stubs could unlock sectoral reallocation benefits. Under composition neutrality, the null result is a structural feature of the labor market that no policy redesign can overcome.

The TWFE comparison offers a methodological lesson. The statistically significant healthcare result under TWFE is exactly the kind of finding that would have been reported and interpreted as evidence of sectoral reallocation prior to the recent advances in staggered DiD methodology (Callaway and Sant'Anna, 2021; Goodman-Bacon, 2021; Sun and Abraham, 2021; Baker et al., 2022). That it vanishes under a heterogeneity-robust estimator illustrates both the practical importance of these methods and the potential for existing published results in the EITC literature to be revisited with modern tools.

Finally, the null result should be placed in quantitative context. The 95% confidence interval for the healthcare share effect (−0.89 to +1.27 pp) rules out reallocation effects larger than 5.6% of the baseline share. This is a meaningful bound: we can reject the hypothesis that the EITC causes large sectoral shifts, even if small effects below our detection threshold remain possible. Future work using linked employer-employee data (e.g., the full LEHD) could achieve greater precision by exploiting within-state, within-industry variation and tracking individual workers' sector transitions.

7. Conclusion

The EITC is the dog that didn't bark—at least along the industry dimension. Despite decades of evidence that the credit powerfully expands women's labor force participation, this paper demonstrates that it does not reshape where women work. Low-education women enter the labor force at the same sectoral proportions they would have joined otherwise: roughly one in four in healthcare, one in six in retail, one in eight in food services.

This null is not a failure of the data or the method. The Quarterly Workforce Indicators provide administrative-quality industry classification for the exact demographic group most affected by the EITC. The staggered adoption of state supplements across 27 states generates credible identifying variation. And the Callaway–Sant'Anna estimator eliminates the spurious positive effect that conventional TWFE produces through heterogeneous treatment timing.

The finding suggests that the EITC's poverty-reduction capacity operates entirely through the transfer itself, not through induced earnings mobility across sectors. For policymakers seeking to promote upward industry mobility among low-income women, the EITC alone is insufficient. Complementary investments in sectoral training, credential recognition, and employer partnerships may be necessary to break the constraint that binds workers to their default sector. The credit opens the door to employment; further policy is needed to direct workers through it.

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," in Timothy Dunne, J. Bradford Jensen, and Mark J. Roberts, eds., *Producer Dynamics: New Evidence from Micro Data*, University of Chicago Press, 2009, pp. 149–230.

Abraham, Katharine G., John Haltiwanger, Kristin Sandusky, and James R. Spletzer, "Exploring Differences in Household vs. Establishment Measures of Employment," *Journal of Labor Economics*, 2013, *31* (2), S129–S172.

Autor, David H., "Why Are There Still So Many Jobs? The History and Future of Workplace Automation," *Journal of Economic Perspectives*, 2015, *29* (3), 3–30.

_ **, Frank Levy, and Richard J. Murnane**, "The Skill Content of Recent Technological Change: An Empirical Exploration," *Quarterly Journal of Economics*, 2003, *118* (4), 1279–1333.

Baker, Andrew C., David F. Larcker, and Charles C.Y. Wang, "How Much Should We Trust Staggered Difference-in-Differences Estimates?," *Journal of Financial Economics*, 2022, *144* (2), 370–395.

Bastian, Jacob E. and Maggie R. Jones, "Do EITC Expansions Pay for Themselves? Effects on Tax Revenue and Government Transfers," *Journal of Public Economics*, 2021, *196*, 104355.

Bayer, Patrick, Stephen L. Ross, and Giorgio Topa, "Place of Work and Place of Residence: Informal Hiring Networks and Labor Market Outcomes," *Journal of Political Economy*, 2008, *116* (6), 1150–1196.

Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan, "How Much Should We Trust Differences-in-Differences Estimates?," *Quarterly Journal of Economics*, 2004, *119* (1), 249–275.

Borusyak, Kirill, Xavier Jaravel, and Jann Spiess, "Revisiting Event-Study Designs: Robust and Efficient Estimation," *Review of Economic Studies*, 2024, *91* (6), 3253–3285.

Bound, John, Charles Brown, and Nancy Mathiowetz, "Measurement Error in Survey Data," in James J. Heckman and Edward Leamer, eds., *Handbook of Econometrics*, Vol. 5, Elsevier, 2001, pp. 3705–3843.

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

Chetty, Raj, Adam Looney, and Kory Kroft, "Salience and Taxation: Theory and Evidence," *American Economic Review*, 2009, *99* (4), 1145–1177.

_ , John N. Friedman, and Emmanuel Saez, "Using Differences in Knowledge Across Neighborhoods to Uncover the Impacts of the EITC on Earnings," *American Economic Review*, 2013, *103* (7), 2683–2721.

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

Dube, Arindrajit, "Minimum Wages and the Distribution of Family Incomes," *American Economic Journal: Applied Economics*, 2019, *11* (4), 268–304.

Eissa, Nada and Hilary Williamson Hoynes, "Behavioral Responses to Taxes: Lessons from the EITC and Labor Supply," *Tax Policy and the Economy*, 2006, *20*, 73–110.

_ and Jeffrey B. Liebman, "Labor Supply Response to the Earned Income Tax Credit," *Quarterly Journal of Economics*, 1996, *111* (2), 605–637.

Goodman-Bacon, Andrew, "Difference-in-Differences with Variation in Treatment Timing," *Econometrica*, 2021, *89* (5), 2291–2333.

Halpern-Meekin, Sarah, Kathryn Edin, Laura Tach, and Jennifer Sykes, *It's Not Like I'm Poor: How Working Families Make Ends Meet in a Post-Welfare World*, University of California Press, 2015.

Hotz, V. Joseph and John Karl Scholz, "Examining the Effect of the Earned Income Tax Credit on the Labor Market Participation of Families on Welfare," Working Paper 11968, National Bureau of Economic Research 2006.

Internal Revenue Service, "Earned Income and Earned Income Tax Credit (EITC) Tables," Publication 596, Internal Revenue Service, Washington, DC 2023.

Ioannides, Yannis M. and Linda Datcher Loury, "Job Information Networks, Neighborhood Effects, and Inequality," *Journal of Economic Literature*, 2004, *42* (4), 1056–1093.

Leigh, Andrew, "Who Benefits from the Earned Income Tax Credit? Incidence among Recipients, Coworkers and Firms," *B.E. Journal of Economic Analysis & Policy*, 2010, *10* (1), Article 45.

Meyer, Bruce D. and Dan T. Rosenbaum, "Welfare, the Earned Income Tax Credit, and the Labor Supply of Single Mothers," *Quarterly Journal of Economics*, 2001, *116* (3), 1063–1114.

Nichols, Austin and Jesse Rothstein, "The Earned Income Tax Credit," in Robert A. Moffitt, ed., *Economics of Means-Tested Transfer Programs in the United States*, Vol. 1, University of Chicago Press, 2016, pp. 137–218.

Pierce, Justin R. and Peter K. Schott, "The Surprisingly Swift Decline of US Manufacturing Employment," *American Economic Review*, 2016, *106* (7), 1632–1662.

Romich, Jennifer L. and Thomas Weisner, "How Families View and Use the EITC: Advance Payment versus Lump Sum Delivery," *National Tax Journal*, 2006, *53* (4), 1245–1265. Part 2.

Ross, Stephen L. and Yves Zenou, "Are Low-Wage Jobs Bad Jobs? Empirical Evidence on Job Quality," 2017. Working Paper.

Rothstein, Jesse, "Is the EITC as Good as an NIT? Conditional Cash Transfers and Tax Incidence," *American Economic Journal: Economic Policy*, 2010, *2* (1), 177–208.

Saez, Emmanuel, "Do Taxpayers Bunch at Kink Points?," *American Economic Journal: Economic Policy*, 2010, *2* (3), 180–212.

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

Sykes, Jennifer, Kathryn Edin, and Sarah Halpern-Meekin, "Dignity and Dreams: What the Earned Income Tax Credit (EITC) Means to Low-Income Families," *American Sociological Review*, 2015, *80* (2), 243–267.

Williams, Erica and Samantha Waxman, "State Earned Income Tax Credits and Minimum Wages Work Best Together," Report, Center on Budget and Policy Priorities 2019.

A. Data Appendix

A.1 State EITC Adoption Timeline

The following states adopted state-level EITC supplements between 1984 and 2018 and are coded as treated in the analysis: Wisconsin (1984), Maryland (1987), Vermont (1988), Iowa (1990), New York (1994), Massachusetts and Oregon (1997), Kansas (1998), Colorado and Indiana (1999), District of Columbia, Illinois, Maine, and New Jersey (2000), Rhode Island (2001), Oklahoma (2002), Delaware, Nebraska, and Virginia (2006), New Mexico (2007), Louisiana and Michigan (2008), Connecticut (2011), Ohio (2013), California (2015), Hawaii and South Carolina (2018). North Carolina adopted a state EITC in 2008 but repealed it in 2014 and is excluded from the treatment group.

The 24 never-treated states are: Alabama, Alaska, Arizona, Arkansas, Florida, Georgia, Idaho, Kentucky, Minnesota, Mississippi, Missouri, Montana, Nevada, New Hampshire, North Carolina, North Dakota, Pennsylvania, South Dakota, Tennessee, Texas, Utah, Washington, West Virginia, and Wyoming.

A.2 QWI Variable Definitions

Employment (`Emp`) is the count of jobs held on the first day of the quarter, averaged across four quarters to produce annual measures. Education is coded in four groups: E1 (less than high school), E2 (high school or equivalent), E3 (some college or associate degree), and E4+E5 (bachelor's degree or above). Industry is classified at the 2-digit NAICS level. I aggregate E1 and E2 as "low education" and E4+E5 as "high education" for the triple-difference specification.

A.3 Industry Classification

The five focal industries and their NAICS codes are: Healthcare and Social Assistance (62), Accommodation and Food Services (72), Retail Trade (44-45), Administrative and Support and Waste Management and Remediation Services (56), and Manufacturing (31-33). Together these sectors account for approximately 68% of low-education women's employment on average across states and years.

B. Identification Appendix

B.1 Pre-Trends Assessment

The event study in Table 3 provides the formal pre-trends test. For food services, all pre-treatment coefficients ($t = -6$ through $t = -2$) are small and statistically insignificant, with magnitudes below 0.3 percentage points. For healthcare, the pre-treatment coefficients at $t = -5$ and $t = -6$ show negative point estimates of approximately 1.2–1.5 percentage points with marginal significance. These long-horizon estimates aggregate across fewer treatment cohorts (only states adopting after 2007 contribute to $t = -6$) and may reflect idiosyncratic pre-treatment variation rather than systematic trend differences. The coefficients at $t = -2$ through $t = -4$ are small and insignificant, supporting parallel trends in the years immediately preceding treatment.

B.2 Control Group Validation

The 24 never-treated states serve as the comparison group for all Callaway–Sant'Anna estimates. These states span diverse regions (Southern, Mountain West, Plains) and economic structures, mitigating concerns that never-treated states are systematically different from EITC-adopting states. The key requirement is not that treated and control states are identical in levels, but that their industry share *trends* would have been parallel absent treatment.

C. Robustness Appendix

The robustness checks in Table 5 demonstrate stability of the TWFE healthcare coefficient across subsamples. Restricting to post-2001 adopters (13 cohorts) yields $\beta = 0.012$, confirming that early adopters are not driving the result. Excluding the Great Recession (2008–2010) produces $\beta = 0.013$, ruling out recession-specific composition effects. Splitting by refundability shows comparable effects for both credit types.

The manufacturing placebo ($\beta = 0.002$, SE $= 0.004$) provides a useful benchmark: the EITC does not predict changes in a sector where employment is driven by trade and technology rather than low-wage labor supply.

D. Standardized Effect Sizes

Table 6: Standardized Effect Sizes for Main Outcomes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Healthcare share	0.0019	0.0055	0.0478	0.039	0.115	Small positive
Food services share	0.0053	0.0040	0.0424	0.126	0.094	Moderate positive
Retail share	-0.0007	0.0021	0.0281	-0.026	0.073	Small negative
Log total employment	0.0536	0.0913	1.0106	0.053	0.090	Moderate positive

Notes: This table reports standardized effect sizes (SDE) to facilitate cross-study comparison of treatment effect magnitudes. For binary (0/1) treatment, SDE = $\hat{\beta}/\text{SD}(Y)$. SD(Y) is the unconditional standard deviation of the outcome variable from Table 1.

Research question: Does state EITC adoption change the industry composition of low-education women's employment? **Treatment:** Binary — state adopted a state-level EITC supplement. **Data:** Census QWI, 2001–2023, state×year panel. **Method:** Callaway–Sant'Anna (2021) staggered DiD with "never treated" control group, state-clustered SEs. **Sample:** 1147 state-year observations across 51 states.

Classification thresholds: large negative (< -0.15), moderate negative (-0.15 to -0.05), small negative (-0.05 to -0.005), null (-0.005 to 0.005), small positive (0.005 to 0.05), moderate positive (0.05 to 0.15), large positive (> 0.15). Classification labels refer to the magnitude of the standardized point estimate, not to statistical significance. "Null" denotes a near-zero effect size ($|\text{SDE}| < 0.005$), not a failure to reject a null hypothesis.