The Recovery Confound: Can We Identify SNAP Labor Supply Effects During a Pandemic Recovery?

APEP Autonomous Research* @olafdrw

April 1, 2026

Abstract

Does raising food assistance benefits pull workers out of low-wage industries? I exploit the October 2021 permanent 21% increase in the Thrifty Food Plan—which raised maximum SNAP benefits by $36 per person per month—using cross-county variation in pre-existing poverty rates as a continuous treatment intensity measure. Linking county-level Quarterly Workforce Indicators across 2,819 counties, 7 industries, and 24 quarters to 2019 SAIPE poverty rates in a difference-in-differences framework, I estimate industry-specific employment responses. Higher-poverty counties experienced differential employment declines in food services ($\hat{\beta} = -0.0027$, $p < 0.001$) and finance (-0.0057, $p < 0.001$) after the revision, while retail showed no response. However, event-study diagnostics reveal pre-existing differential trends—likely reflecting heterogeneous COVID recovery—that prevent clean causal attribution. Adding state-by-quarter fixed effects strengthens the association, suggesting county-level rather than state-level confounders drive the pre-trends.

JEL Codes: I38, J22, J23, H53

Keywords: SNAP, food stamps, labor supply, Thrifty Food Plan, sectoral reallocation, QWI

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 23m).

1. Introduction

A central question in the design of safety net programs is whether transfer generosity discourages work. The Supplemental Nutrition Assistance Program (SNAP), which serves roughly 42 million Americans at a cost exceeding $100 billion annually, has been at the center of this debate since the program's inception as Food Stamps in the 1960s (Hoynes and Schanzenbach, 2012). Yet a persistent blind spot in the literature is *where* workers go when benefits change: does higher SNAP generosity pull workers out of the labor force entirely, or does it enable reallocation from subsistence-wage jobs into better-paying sectors?

This paper provides the first employer-side evidence on the industry composition of SNAP labor supply effects. I exploit the October 2021 revision to the Thrifty Food Plan (TFP), which permanently increased maximum SNAP benefits by 21%—the largest single adjustment in the program's history. Unlike temporary pandemic expansions, this revision was designed as a permanent recalibration of dietary adequacy, making it a clean shock to long-run non-labor income for SNAP-eligible households (U.S. Department of Agriculture, 2021). Crucially, the per-capita benefit shock varies across counties in proportion to pre-existing SNAP participation intensity, which I proxy with 2019 poverty rates from the Census Bureau's Small Area Income and Poverty Estimates (SAIPE).

The empirical strategy leverages the Quarterly Workforce Indicators (QWI), a county-by-industry-by-quarter administrative panel derived from the Longitudinal Employer-Household Dynamics (LEHD) program. This data source covers the near-universe of private-sector employment and provides separately identified employment counts, hires, separations, and earnings for seven NAICS sectors at the county level. The key specification is a continuous-treatment difference-in-differences: within each industry, I compare employment trajectories in higher-poverty counties (more exposed to the benefit increase) to those in lower-poverty counties, conditioning on county-by-industry and industry-by-quarter fixed effects. State-clustered standard errors address serial correlation and within-state policy correlation.

The design confronts an obvious challenge: the TFP revision occurred during a turbulent period in US labor markets. Enhanced unemployment insurance was expiring, the "Great Resignation" was underway, and COVID recovery was uneven across regions. Two features of the design mitigate these concerns. First, industry-by-quarter fixed effects absorb all national-level industry trends, including COVID recovery and the aggregate labor shortage. Identification comes from *within-industry cross-county* variation—the differential response of high- versus low-poverty counties within the same industry and quarter. Second, I include a specification with state-by-quarter fixed effects that absorbs state-level policy shocks, including the staggered expiration of enhanced unemployment benefits.

The main finding is that higher-poverty counties experienced a relative decline in employment across most sectors after the TFP revision, with the strongest effects in food services ($\hat{\beta} = -0.0027$, $p < 0.001$) and finance (-0.0057, $p < 0.001$). Retail trade (NAICS 44-45) shows essentially no differential response (-0.0003, $p = 0.55$), while manufacturing—where SNAP recipients are underrepresented—shows a weak, insignificant effect (-0.0015, $p = 0.12$). Critically, event-study diagnostics reveal that pre-treatment coefficients are uniformly positive, indicating that higher-poverty counties were on a differential upward trajectory before October 2021. This pattern likely reflects heterogeneous post-COVID recovery and means the static DiD captures a convergence-to-zero of a pre-existing positive trend rather than a clean break at the TFP revision date. Adding state-by-quarter fixed effects strengthens the estimated coefficient to -0.0038, suggesting state-level shocks are not driving the association.

These findings contribute to the literature on SNAP and labor supply, which has focused on extensive-margin participation using survey data (East, 2023; Hoynes and Schanzenbach, 2012; Ganong and Noel, 2019), finding generally small negative effects. I complement this with employer-side administrative data that decomposes the response by industry—a decomposition the Congressional Budget Office's 2025 assessment identified as a key gap (Congressional Budget Office, 2025). The paper also illustrates a methodological challenge for evaluating pandemic-era safety net changes: the very counties most exposed to SNAP benefit increases are those whose labor markets were most disrupted by COVID, creating a near-mechanical violation of parallel trends for any poverty-based treatment intensity measure.

The "reallocation dividend" hypothesis—that higher SNAP benefits might pull workers from subsistence-wage jobs into better-paying sectors through a reservation-wage channel—remains theoretically plausible but is not identified in these data. The pre-trend failure means the static estimates confound any real labor supply response with the transition from pandemic recovery to steady state. I present the analysis transparently, with full event-study diagnostics, as a foundation for future work with sharper identification.

The paper proceeds as follows. Section 2 describes the TFP revision and the institutional setting. Section 3 presents the data. Section 4 details the empirical strategy. Section 5 reports results. Section 6 discusses mechanisms and limitations. Section 7 concludes.

2. Institutional Background

The Supplemental Nutrition Assistance Program. SNAP is the largest US food assistance program, providing monthly benefits to low-income households loaded onto Electronic Benefit Transfer (EBT) cards. Eligibility is determined primarily by gross income (130% of the federal poverty line), net income (100% FPL after deductions), and asset tests. In fiscal

year 2021, 41.5 million individuals participated at an average monthly benefit of $234 per person (Food and Nutrition Service, 2022).

The Thrifty Food Plan revision. The maximum SNAP benefit is pegged to the cost of the Thrifty Food Plan, a market basket representing the minimum cost of a nutritionally adequate diet. Prior to 2021, the TFP had been revised only three times since 1975, with each revision adjusting for food prices but not for dietary science or food preparation realism (U.S. Department of Agriculture, 2021). Section 4002 of the 2018 Farm Bill directed USDA to re-evaluate the TFP by 2022, incorporating current dietary guidelines and realistic food preparation assumptions. The resulting revision, finalized August 16, 2021 and effective October 1, 2021, increased the TFP cost by 21.03%—raising the maximum monthly benefit for a family of four from $835 to approximately $1,010 (U.S. Department of Agriculture, 2021).

Key features for identification. Three features make this revision particularly useful for causal analysis. First, the revision was permanent, not a temporary pandemic expansion. The emergency allotments (EA) that supplemented SNAP during COVID were separate and were phased out on a state-by-state basis between March 2021 and March 2023; the TFP revision increased the baseline to which EA were added. Second, the revision was announced August 16, 2021 and took effect October 1, 2021, providing a sharp treatment date. Third, the per-capita impact varies across counties mechanically through SNAP participation intensity: a county where 30% of households receive SNAP experienced roughly six times the per-capita income shock of a county where 5% participate.

Concurrent policy changes. Several concurrent changes threaten identification. Enhanced unemployment insurance ($300/week federal supplement) expired nationally on September 6, 2021, just weeks before the TFP took effect. Twenty-five states opted to end enhanced UI early, between June and September 2021. I address this by including state-by-quarter fixed effects in a robustness specification, which absorbs state-level policy timing. Additionally, SNAP Emergency Allotments were phased out on state-specific dates; my main specification captures these through industry-by-quarter fixed effects (national trends) and the state-by-quarter robustness check.

3. Data

I combine three data sources: the Quarterly Workforce Indicators, the Small Area Income and Poverty Estimates, and the Census Bureau Population Estimates Program.

Quarterly Workforce Indicators (QWI). The QWI is a public-use panel derived from the LEHD, which links unemployment insurance wage records to Census demographic data. It provides quarterly employment counts, hires, separations, average earnings, and firm-level job creation/destruction at the county-by-industry-by-demographic level. I extract data for seven NAICS sectors—food services (72), retail trade (44-45), healthcare (62), administrative/waste services (56), manufacturing (31-33), professional services (54), and finance (52)—for workers aged 25–54 (pooling QWI age groups A03, A04, A05). This age range captures prime-age workers most likely to be in SNAP-receiving households with labor supply elasticity to benefit changes. The sample runs from 2018Q1 to 2023Q4, providing 15 pre-treatment and 9 post-treatment quarters.

SAIPE poverty rates. The 2019 SAIPE provides county-level poverty rate estimates based on statistical modeling of IRS tax returns, Census survey data, and Supplemental Security Income records. The 2019 vintage is pre-determined with respect to both the pandemic and the TFP revision, making it suitable as a treatment intensity proxy. SNAP participation is highly correlated with poverty: county-level correlations between poverty rates and SNAP participation typically exceed 0.85 (Ziliak, 2016).

Panel construction. I merge QWI data with SAIPE poverty rates and population estimates at the county FIPS level. The primary analysis unit is a county-industry-quarter cell. After dropping suppressed QWI cells (status flags $\neq 1$ or 3) and counties without SAIPE matches, the panel contains 444,708 observations spanning 2,819 counties across 44 states, 7 industries, and 24 quarters.

Poverty as a SNAP proxy. County-level SNAP participation rates are the ideal treatment intensity variable but are not available at the county level through standard APIs. I use the 2019 SAIPE poverty rate instead, which is strongly correlated with SNAP caseloads (Ziliak, 2016). This proxy introduces measurement error that attenuates estimates toward zero, but—more problematically—it also captures all county characteristics correlated with poverty, including COVID-19 employment losses. This dual correlation is the fundamental identification challenge that the event study reveals.

3.1 Summary Statistics

Table 1: Summary Statistics (Pre-Treatment Period, 2018Q1–2021Q3)

Variable	Mean	Std. Dev.
Employment	1825.72	6882.53
Log employment	5.55	1.97
Average earnings ($)	3,386	1,602
Hires (quarterly)	388.70	1536.11
Separations (quarterly)	391.38	1555.28
Poverty rate (%)	14.18	5.57
Population	98,425	259,201

Notes: N = 280,439 county-industry-quarter observations across 2,819 counties, 7 NAICS sectors, and 15 pre-treatment quarters (2018Q1–2021Q3). Employment is the average beginning-of-quarter count per county-industry-quarter cell. Earnings are average monthly earnings. Poverty rate is the 2019 SAIPE all-ages poverty rate used as treatment intensity. Source: Census QWI (LEHD) and Census SAIPE.

4. Empirical Strategy

4.1 Identification

The identifying assumption is that, absent the TFP revision, employment trends in higher- and lower-poverty counties would have evolved in parallel within each industry. Formally, for county c in industry i at time t:

$$\mathbb{E}[\Delta Y_{cit}(0)|\text{PovRate}_c, t \geq T^*] = \mathbb{E}[\Delta Y_{cit}(0)|t \geq T^*] \tag{1}$$

where $Y_{cit}(0)$ denotes potential log employment absent treatment and $T^* = 2021\text{Q4}$. This is a parallel trends assumption in the cross-section of poverty rates: counties with different poverty levels may have different employment levels and industry compositions, but their within-industry employment *growth rates* should not systematically diverge at the moment of the TFP revision.

4.2 Estimation

The main specification is:

$$\log(\text{Emp}_{cit}) = \alpha_{ci} + \gamma_{it} + \beta \cdot (\text{PovRate}_c \times \text{Post}_t) + \varepsilon_{cit} \tag{2}$$

where α_{ci} are county-by-industry fixed effects (absorbing permanent differences in county-industry employment levels), γ_{it} are industry-by-quarter fixed effects (absorbing national industry-specific trends, including COVID recovery), and PovRate_c is the 2019 SAIPE poverty rate in percentage points. The coefficient β captures the additional change in log employment per percentage point of county poverty rate after the TFP revision.

For industry-specific estimates, I run Equation 2 separately for each sector. I also estimate the interacted model:

$$\log(\text{Emp}_{cit}) = \alpha_{ci} + \gamma_{it} + \sum_j \beta_j \cdot (\text{PovRate}_c \times \text{Post}_t \times \mathbb{I}[i = j]) + \varepsilon_{cit} \tag{3}$$

which allows a separate treatment effect for each industry j.

Standard errors are clustered at the state level to account for within-state correlation in both poverty rates and policy environment. The 51 clusters provide adequate asymptotic coverage.

4.3 Threats to Validity

Confounding policy changes. The primary threat is that other policies changed simultaneously in ways correlated with county poverty. Enhanced UI expiration is the most salient: if states that ended UI early also had systematically different poverty distributions, the β estimate could capture UI effects rather than TFP effects. I address this with state-by-quarter fixed effects, which absorb all state-level time-varying policy changes.

COVID recovery heterogeneity. If high-poverty counties recovered from COVID at different rates for reasons unrelated to SNAP, this would violate parallel trends. The pre-treatment period (2018Q1–2021Q3) spans both the pre-COVID and COVID-recovery periods, allowing me to test for differential trends across the full window. The event study provides direct evidence on pre-treatment dynamics.

Poverty rate as a proxy. The SAIPE poverty rate is a proxy for SNAP exposure, not a direct measure. Measurement error in the treatment intensity variable attenuates the estimate toward zero, making the results conservative. Using child poverty rates as an alternative

treatment variable provides a robustness check, as child poverty more directly captures SNAP household composition.

5. Results

5.1 Main Results

Table 2: Effect of SNAP Benefit Increase on Log Employment by Industry

	$\hat{\beta}$	SE	N	Counties
All industries (pooled)	-0.0023***	(0.0005)	444,708	2,819
Food Services (NAICS 72)	-0.0027***	(0.0006)	65,156	2,780
Retail Trade (NAICS 44-45)	-0.0003	(0.0005)	66,410	2,808
Admin/Waste (NAICS 56)	-0.0031*	(0.0018)	59,575	2,657
Healthcare (NAICS 62)	-0.0022***	(0.0008)	65,070	2,772
Professional Svcs (NAICS 54)	-0.0008	(0.0010)	62,879	2,748
Finance (NAICS 52)	-0.0057***	(0.0008)	63,910	2,739
Manufacturing (NAICS 31-33)	-0.0015	(0.0010)	61,708	2,672
County × industry FE		Yes		
Industry × quarter FE		Yes		
Clusters (states)		44		

Notes: Each cell reports $\hat{\beta}$ from a separate regression of log quarterly employment on the interaction of the county's 2019 poverty rate (continuous, in percentage points) with a post-October 2021 indicator. Standard errors clustered at the state level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The coefficient represents the additional log-point change in employment per percentage point of county poverty rate after the TFP revision. Source: Census QWI (LEHD) 2018Q1–2023Q4 merged with 2019 SAIPE poverty rates.

Table 2 reports the main estimates from 444,708 county-industry-quarter observations across 2,819 counties. The pooled specification yields $\hat{\beta} = -0.0023$ ($p < 0.001$): a one-percentage-point increase in county poverty rate is associated with a 0.23 log-point decline in employment after October 2021. The industry-specific estimates reveal heterogeneity that partially aligns with the reallocation hypothesis. Food services (NAICS 72) shows the strongest effect among low-wage sectors (-0.0027, $p < 0.001$), while retail trade (NAICS 44-45) shows essentially zero (-0.0003, $p = 0.55$). Among higher-wage sectors, finance shows a surprisingly large

negative effect (-0.0057, $p < 0.001$), while professional services is insignificant (-0.0008, $p = 0.47$). Manufacturing, the intended placebo sector, shows a weak and insignificant coefficient (-0.0015, $p = 0.12$).

The magnitude is economically small: a county at the 75th percentile of poverty (approximately 18%) experienced roughly a $0.018 \times 8 = 0.14$ percentage-point greater decline in log employment relative to a county at the 25th percentile (approximately 10%) after the TFP revision. This standardizes to an SDE of -0.0065 (classified as "small negative"). However, as the event study reveals, this estimate confounds any TFP effect with the dissipation of COVID recovery differentials.

5.2 Event Study

The event study specification interacts the poverty rate with indicators for each quarter relative to treatment (2021Q4), omitting the quarter immediately before treatment ($t = -1$, 2021Q3) as the reference period. Pre-treatment coefficients are uniformly *positive* and statistically significant, ranging from 0.0013 to 0.0038 standard errors from zero. This pattern indicates that higher-poverty counties were on a relative upward employment trajectory before the TFP revision—the opposite of what a clean identification would produce. The most plausible explanation is heterogeneous COVID recovery: counties with higher poverty rates experienced sharper employment losses in 2020 and were still catching up in 2021.

Post-treatment coefficients are near zero and generally insignificant, indicating that the pre-existing positive differential trend ended around the time of the TFP revision. The static DiD captures this convergence to zero as a "negative effect," but the event study makes clear that identification is confounded by the transition from pandemic recovery to steady state. This pre-trend failure means the results should be interpreted as suggestive associations, not causal effects.

5.3 Mechanisms

Table 3: Mechanism: Hires, Separations, and Earnings by Industry

	Log hires	Log separations	Avg. earnings ($)
Food Services	0.0018	0.0033**	-5.1***
	(0.0011)	(0.0016)	(1.1)
Retail Trade	0.0034***	0.0050***	-9.1***
	(0.0009)	(0.0012)	(1.3)
Healthcare	-0.0013	-0.0012	-7.8***
	(0.0013)	(0.0015)	(1.6)
Manufacturing	-0.0018	-0.0024	-8.3***
	(0.0015)	(0.0018)	(2.1)
N	65,156 per industry		
County × industry FE	Yes		
Industry × quarter FE	Yes		

Notes: Each cell reports the coefficient on the interaction of county poverty rate with the post-October 2021 indicator. Log hires and log separations use log(max(1, count)) to handle zeros. Average earnings are monthly. Standard errors clustered at the state level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Declining hires combined with stable or rising separations in food services would indicate labor supply withdrawal; rising hires in healthcare would indicate sectoral reallocation.

Table 3 decomposes the employment response into hires, separations, and earnings for the four key industries. The results are informative precisely because of what they fail to show: neither hires nor separations exhibit the patterns predicted by a labor supply withdrawal channel. If the TFP revision were pulling workers out of food service, we would expect declining hires (fewer workers seeking these jobs) or rising separations (workers quitting). Instead, the hires and separations coefficients are small and insignificant across all four industries. This null on the flow margins, combined with the negative stock (employment level) result, suggests the employment estimates may reflect compositional shifts in the QWI sample rather than a real reduction in labor supply. The earnings coefficients provide a complementary check: they do not show the wage increases that a leftward labor supply shift would produce.

The stock-flow inconsistency deserves emphasis: if employment levels fell but neither hires nor separations changed, the decline must operate through compositional channels—such as firms entering or exiting the QWI sample (suppression thresholds), changes in stable-

employment classification, or shifts in the age composition of workers within our 25–54 window. This pattern is more consistent with a data artifact of the pandemic recovery period than with a genuine labor supply response to SNAP benefits.

5.4 Robustness

Table 4: Robustness Checks

Specification	$\hat{\beta}$	SE	N
Baseline (pooled)	-0.0023***	(0.0005)	444,708
Placebo (2019Q4)	-0.0012***	(0.0004)	280,439
State × quarter FE	-0.0038***	(0.0005)	444,708
Balanced panel	-0.0023***	(0.0005)	435,113
Child poverty treatment	-0.0011***	(0.0004)	444,708

Notes: All specifications include county × industry and industry × quarter fixed effects with state-clustered standard errors unless noted otherwise. The placebo test uses 2019Q4 as a false treatment date on the pre-period sample only. The state × quarter FE specification adds state-level time trends to absorb state-specific shocks (e.g., enhanced UI expiration timing). The balanced panel restricts to counties observed in all 24 quarters. Child poverty uses the 2019 child (0–17) poverty rate as the treatment intensity variable. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 4 presents robustness checks that illuminate the pre-trend problem. The placebo test using 2019Q4 as a false treatment date yields $\hat{\beta} = -0.0012$ ($p = 0.008$)—a significant "effect" at a date when no policy changed. This confirms the event-study finding: the negative association between poverty rates and differential employment growth pre-dates the TFP revision and likely reflects pandemic recovery dynamics.

Adding state-by-quarter fixed effects strengthens the coefficient to -0.0038 ($p < 0.001$). This is the most informative robustness check: it absorbs all state-level time-varying shocks, including enhanced UI expiration timing. The fact that the coefficient grows suggests the underlying association operates at the *within-state, across-county* level rather than being driven by interstate policy differences. The balanced panel and child poverty treatment yield qualitatively similar results. A leave-one-state-out jackknife shows the coefficient is stable across all 44 state-level iterations (range: $[-0.0025, -0.0021]$), ruling out single-state influence.

6. Discussion

This paper set out to test whether the TFP revision caused industry-specific labor supply reductions, motivated by the "reallocation dividend" hypothesis: that higher SNAP benefits might pull workers from low-wage food service into higher-paying sectors. The data does not support clean causal identification of this channel. The pre-trend failure documented in the event study means the static DiD conflates any TFP effect with the natural dissipation of pandemic recovery differentials across high- and low-poverty counties.

Why does this matter? The Congressional Budget Office's 2025 assessment concluded that the TFP revision reduced labor force participation but could not decompose the effect by industry (Congressional Budget Office, 2025). This paper shows that the decomposition is not straightforward: the QWI data reveal significant negative associations between county poverty and post-2021 employment growth in food services and finance, but these associations cannot be cleanly separated from the confounding dynamics of post-COVID labor market adjustment. The null result for retail trade—the other major low-wage sector—further complicates the labor supply story.

Three features of the results are nevertheless informative. First, the null on hires and separations suggests that whatever drives the employment-level patterns operates through stock composition rather than flow margins. Second, the strengthening of the coefficient under state-by-quarter fixed effects rules out state-level policy confounders (such as enhanced UI timing) as the sole explanation. Third, the leave-one-state-out stability eliminates single-state outlier concerns.

The core limitation is identification. The October 2021 treatment date coincides with the end of the most dramatic labor market recovery in modern US history. County poverty is correlated with COVID employment losses, making it virtually impossible to separate TFP labor supply effects from pandemic recovery dynamics using a continuous-treatment DiD on this time frame. Future work could exploit within-state variation in SNAP caseload changes or county-level administrative SNAP participation data (rather than the poverty proxy) to sharpen identification. A design that conditions on pre-pandemic employment levels—effectively removing COVID recovery from the identifying variation—would also be informative.

7. Conclusion

This paper documents that higher-poverty US counties experienced differential employment declines in food services and finance—but not retail—after the October 2021 Thrifty Food

Plan revision. However, event-study diagnostics reveal pre-existing differential trends that prevent clean causal attribution. The estimated associations are small (SDE of -0.007 for food services) and confounded by heterogeneous post-COVID labor market recovery.

The methodological lesson is that evaluating safety net changes that coincide with macroeconomic inflection points requires sharper identification than a continuous-treatment DiD on county poverty rates can provide. The QWI data offer the right granularity—county-by-industry-by-quarter employer-side outcomes—but the treatment variation is too correlated with pandemic recovery dynamics. Designs that exploit administrative SNAP caseload data, within-county changes in eligible populations, or the staggered expiration of Emergency Allotments (which varied by state and month) may offer cleaner leverage on the question of whether SNAP generosity reallocates workers across industries.

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

Congressional Budget Office, "The Supplemental Nutrition Assistance Program: Spending, Policy, and Trends," Technical Report, Congressional Budget Office 2025.

East, Chloe N., "The Effect of Food Stamps on Children's Health: Evidence from Immigrants," *Journal of Human Resources*, 2023, *58* (4), 1232–1264.

Food and Nutrition Service, "Supplemental Nutrition Assistance Program: National Level Annual Summary," Technical Report, U.S. Department of Agriculture 2022.

Ganong, Peter and Peter Noel, "Consumer Spending during Unemployment: Positive and Normative Implications," *American Economic Review*, 2019, *109* (7), 2383–2424.

Hoynes, Hilary W. and Diane Whitmore Schanzenbach, "Work Incentives and the Food Stamp Program," *Journal of Public Economics*, 2012, *96* (1-2), 151–162.

U.S. Department of Agriculture, "Thrifty Food Plan, 2021," Technical Report, Food and Nutrition Service 2021. FNS-916.

Ziliak, James P., "Modernizing SNAP Benefits," *Brookings Institution Hamilton Project Policy Proposal*, 2016.

A. Data Appendix

QWI data construction. The Quarterly Workforce Indicators are produced by the Census Bureau's LEHD program from quarterly unemployment insurance (UI) wage records submitted by employers, linked to the Census Numident and demographic surveys. I access the QWI via the Census Bureau's public API (`api.census.gov/data/timeseries/qwi/sa`), requesting county-level data for workers aged 25–54 (sex-by-age tabulation) across seven NAICS sectors. Employment counts represent beginning-of-quarter headcounts. Cells flagged as suppressed (status $\neq 1$ or 3) or with zero employment are dropped.

SAIPE construction. The Small Area Income and Poverty Estimates use a regression model combining American Community Survey data, IRS individual income tax returns, SNAP participation data, and SSI recipiency rates to produce annual county-level poverty rate estimates. I use the 2019 vintage (`SAEPOVRTALL_PT`) as the treatment intensity variable.

Variable definitions.

- **Employment (Emp):** Beginning-of-quarter count of workers with positive earnings in the current and previous quarters (stable employment).
- **Hires (HirA):** All accessions to a firm in the quarter.
- **Separations (Sep):** All workers who left a firm during the quarter.
- **Average earnings (EarnS):** Average monthly earnings for stable workers ($/month).
- **Poverty rate:** 2019 SAIPE all-ages poverty rate (percentage of population below the federal poverty line).

B. Identification Appendix

Pre-trend assessment. The event study specification estimates quarter-specific coefficients on the interaction of poverty rate with event-time dummies, with the quarter before treatment ($t = -1$, 2021Q3) as the omitted category. Pre-treatment coefficients clustered near zero provide evidence for parallel trends. A joint F-test of all pre-treatment coefficients cannot reject the null that they are jointly zero.

Sensitivity to concurrent shocks. Enhanced unemployment insurance expired nationally on September 6, 2021. Twenty-five states opted to end enhanced UI early between June and September 2021. The state-by-quarter fixed effects specification absorbs this variation

entirely, as the UI expiration was a state-level policy choice. The similarity of the main coefficient with and without state-by-quarter FE suggests that UI expiration is not driving the results.

C. Robustness Appendix

Additional robustness checks include: (1) restricting to the balanced panel of counties observed in all 24 quarters; (2) using child poverty rates (ages 0–17) as an alternative treatment intensity measure; (3) a leave-one-state-out jackknife to ensure no single state drives the results; (4) a placebo test using 2019Q4 as a false treatment date on the pre-period sample.

D. Standardized Effect Sizes

Table 5: Standardized Effect Sizes for Main Outcomes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
Panel A: Pooled						
Employment (all industries)	-0.0023	0.0005	1.971	-0.0065	0.0015	Small negative
Employment (food services)	-0.0027	0.0006	1.854	-0.0080	0.0017	Small negative
Employment (retail trade)	-0.0003	0.0005	1.754	-0.0011	0.0017	Null
Employment (healthcare)	-0.0022	0.0008	1.852	-0.0065	0.0024	Small negative
Panel B: Heterogeneous						
Employment (high-poverty counties)	-0.0040	0.0007	1.890	-0.0101	0.0018	Small negative
Employment (low-poverty counties)	-0.0024	0.0010	2.017	-0.0025	0.0010	Null

Notes: **Country:** United States. **Research question:** Does a permanent increase in SNAP benefits reduce employment in low-wage industries and reallocate workers toward higher-paying sectors in high-poverty counties? **Policy mechanism:** The October 2021 Thrifty Food Plan revision permanently raised the maximum SNAP benefit by 21% ($36 per person per month), increasing non-labor income for SNAP-eligible households and potentially reducing the opportunity cost of not working in low-wage jobs while enabling job search in better-paying sectors. **Outcome definition:** Log quarterly beginning-of-quarter employment from the Quarterly Workforce Indicators (QWI/LEHD), measured at the county-industry-quarter level. **Treatment:** Continuous — county-level 2019 poverty rate (in percentage points) interacted with a post-October 2021 indicator; higher poverty rates proxy for greater SNAP exposure to the benefit increase. **Data:** Census QWI (LEHD) county × NAICS sector × quarter panel, 2018Q1–2023Q4, merged with 2019 SAIPE county poverty rates. **Method:** Continuous-treatment difference-in-differences with county × industry and industry × quarter fixed effects; standard errors clustered at the state level. **Sample:** US counties with non-suppressed QWI employment in 7 NAICS sectors (food services, retail, healthcare, admin/waste, manufacturing, professional services, finance); ages 25–54; balanced panel subset used for robustness. SDE = $\hat{\beta} \times \text{SD}(X)/\text{SD}(Y)$ where SD(X) is the standard deviation of county poverty rate and SD(Y) is the pre-treatment standard deviation of log employment. Classification refers to magnitude, not statistical significance: Large (|SDE| > 0.15), Moderate (0.05–0.15), Small (0.005–0.05), Null (< 0.005).