

The SNAP Bridge: Transitional Benefits and the Welfare Cliff

APEP Autonomous Research* @SocialCatalystLab

March 23, 2026

Abstract

When families leave TANF cash assistance, they must typically reapply separately for SNAP—creating a bureaucratic gap in food assistance. Since 2001, 24 states have adopted transitional SNAP benefits, providing automatic 5-month SNAP continuation for families exiting welfare. I exploit the staggered adoption of these policies to estimate the effect on SNAP participation using Census ACS data and the [Callaway and Sant’Anna \(2021\)](#) estimator. Point estimates suggest a 0.5 percentage point increase in the state SNAP participation rate, but the effect is statistically imprecise ($SE = 0.6$ pp, $p > 0.10$). This imprecision likely reflects signal dilution: TANF exiters represent less than 3% of SNAP households, making aggregate state-level data a blunt instrument for detecting effects at the TANF-exit margin. The consistently positive sign across specifications, combined with flat pre-trends, rules out large negative effects and provides suggestive evidence of a modest positive impact.

JEL Codes: I38, H75, J18

Keywords: SNAP, food stamps, TANF, transitional benefits, welfare cliff, staggered difference-in-differences

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: 1h 16m).

1. Introduction

Millions of American families depend on the safety net not as a single program but as a patchwork of overlapping eligibility windows, application deadlines, and recertification requirements. When a family exits Temporary Assistance for Needy Families (TANF)—whether because they found employment, exhausted time limits, or were sanctioned—they do not automatically retain their Supplemental Nutrition Assistance Program (SNAP) benefits. In most states, they must file a new application, provide documentation, and wait for processing. This bureaucratic gap between programs is a canonical example of the “welfare cliff”: the risk that leaving one program causes a family to lose access to another, not because they are ineligible, but because the administrative machinery fails to connect them (Ziliak, 2015; Currie, 2006).

The consequences of this gap are potentially severe. SNAP is the single largest anti-hunger program in the United States, serving over 40 million people at its peak and providing an average monthly benefit that constitutes a substantial share of food spending for low-income households (Hoynes and Schanzenbach, 2016; Ganong and Liebman, 2018). Losing food assistance during the transition from cash welfare to employment—precisely the moment when families face the highest financial uncertainty—can undermine the goals of welfare reform itself. If families cannot maintain nutrition security while they stabilize in the labor market, the transition from dependence to self-sufficiency becomes harder, not easier (Bitler et al., 2017; East, 2018).

Beginning in 2001, the federal government gave states the option to provide *transitional SNAP benefits*: automatic five-month SNAP continuation for families leaving TANF, with no separate application required. By 2016, twenty-four states and the District of Columbia had adopted this option, creating substantial variation in policy timing that is well-suited to causal identification. Yet despite over two decades of implementation, no study has estimated the causal effect of transitional SNAP benefits on food assistance participation. This paper fills that gap.

I construct a state-year panel of SNAP participation rates from the American Community Survey (ACS) covering 2005–2023, merged with policy adoption dates from the USDA SNAP Policy Database. Using the Callaway and Sant’Anna (2021) estimator for staggered difference-in-differences—which avoids the well-documented biases of two-way fixed effects (TWFE) when treatment effects are heterogeneous (Goodman-Bacon, 2021; Sun and Abraham, 2021; Roth et al., 2023)—I estimate the average treatment effect on the treated (ATT) of transitional SNAP adoption on state-level SNAP participation rates.

The main estimate is a 0.5 percentage point increase in the SNAP participation rate (SE

= 0.6 pp), which is positive but statistically insignificant at conventional levels. The log specification yields a coefficient of 0.037 (SE = 0.074), telling the same story. These estimates are robust across specifications: TWFE produces similar point estimates, the event study shows no evidence of pre-trends, and a randomization inference exercise yields a p -value of 0.69—consistent with a small positive effect swamped by noise.

Why is the effect so imprecise? The answer lies in the denominator. TANF caseloads have declined precipitously since welfare reform, falling from over 4.5 million families in 1996 to roughly 1 million by the mid-2000s (Falk, 2022). In a typical state, the number of families exiting TANF in a given year is a small fraction of total SNAP households. Even if transitional benefits perfectly retained every TANF exiter on SNAP, the effect on the overall state SNAP rate would be modest—on the order of 0.5–2 percentage points in most states. Detecting this signal in aggregate data requires either very long panels or a more targeted outcome.

This paper makes three contributions. First, it provides the first causal estimate of transitional SNAP benefits on food assistance participation, adding to the growing literature on administrative barriers in safety net programs (Currie, 2006; Finkelstein and Notowidigdo, 2019; Homonoff and Somerville, 2021). Even a null result is informative: it tells us that the bureaucratic gap between TANF and SNAP, while real, may be small enough in practice that most families navigate it—or that administrative simplification at the state level is not sufficient to move the aggregate needle.

Second, the paper contributes to the literature on the “welfare cliff” and program interactions. Ziliak (2015) documented the problem of SNAP churn, while Gray et al. (2023) studied the consequences of SNAP emergency allotment expiration during COVID-19. This paper examines a different margin: the cross-program transition from TANF to SNAP, where administrative friction rather than eligibility determines whether families maintain food assistance.

Third, the analysis demonstrates how modern heterogeneity-robust DiD methods perform in settings with potentially small effects. The agreement between Callaway and Sant’Anna (2021) and TWFE estimates is reassuring—it suggests that treatment effect heterogeneity, while a valid concern in staggered adoption designs, does not drive the results here. The paper also implements randomization inference (Fisher, 1935; Heß, 2017) to provide exact finite-sample p -values, appropriate given the modest number of treatment cohorts.

The remainder of the paper proceeds as follows. Section 2 describes the institutional background of transitional SNAP benefits. Section 3 presents the data. Section 4 details the empirical strategy. Section 5 reports results. Section 6 discusses implications and concludes.

2. Institutional Background and Policy Setting

The TANF-SNAP administrative gap. The Supplemental Nutrition Assistance Program (SNAP) and Temporary Assistance for Needy Families (TANF) serve overlapping but distinct populations. SNAP provides monthly food purchasing benefits to households below 130% of the federal poverty line (or 200% in broad-based categorical eligibility states), while TANF provides time-limited cash assistance with work requirements to families with children (Falk, 2022). Most TANF recipients also receive SNAP, but the two programs have separate application processes, eligibility determinations, and recertification cycles.

When a family leaves TANF—whether voluntarily or through sanctions, time limits, or earnings gains—their SNAP case is not automatically adjusted. In states without transitional benefits, the family must file a new SNAP application, provide income verification, and wait for eligibility determination. This process can take 30 days or more, creating a gap in food assistance during a period of acute financial transition (Ziliak, 2015). The problem is compounded by the fact that many TANF exiters are transitioning to low-wage employment and face immediate expenses (transportation, childcare, work clothing) that strain already-tight budgets.

Transitional SNAP benefits. Section 217 of Public Law 106-387 (the FY2001 Agriculture Appropriations Act) authorized states to provide transitional food stamp benefits to families leaving TANF. Under this option, families exiting TANF automatically receive five months of SNAP benefits at their pre-exit benefit level, with no new application or recertification required during the transitional period. At the end of five months, the family must apply through the standard process to continue receiving SNAP.

The policy is designed to serve as a “bridge”—smoothing the administrative transition from TANF to standalone SNAP eligibility. It does not change the benefit amount (which remains based on pre-exit income and household size), nor does it extend the total duration of SNAP eligibility. Rather, it eliminates the application burden during the first five months after TANF exit, reducing the probability that eligible families lose food assistance due to paperwork rather than ineligibility.

State adoption patterns. Adoption of transitional SNAP benefits was staggered across states from 2001 to 2016. Early adopters (2001–2003) included states such as Pennsylvania, Oregon, and Maryland. A second wave of adoption occurred in 2004–2008, and the final adopters implemented the policy between 2009 and 2016. By 2023, twenty-four states and the District of Columbia had adopted the policy. Twenty-seven states (including the remaining jurisdictions) never adopted transitional benefits during the study period.

The staggered timing of adoption creates useful identification variation. States adopted transitional benefits at different points in the business cycle, under different political configurations, and with different TANF caseload trajectories. While adoption is not randomly assigned, the timing variation allows me to control for both state-level time-invariant confounders and national trends using the [Callaway and Sant’Anna \(2021\)](#) estimator.

TANF caseload context. Understanding why the effect on aggregate SNAP participation may be small requires context on TANF caseload sizes. Following the 1996 welfare reform (PRWORA), TANF caseloads declined dramatically—from 4.5 million families in 1996 to roughly 1.7 million by 2005 and under 1 million by 2020 ([Falk, 2022](#)). In a typical state, annual TANF exits might number in the tens of thousands, while total SNAP households number in the hundreds of thousands. Even if transitional benefits retained every single TANF exiter on SNAP (an upper bound), the effect on the state participation rate would be diluted by the much larger non-TANF SNAP population.

This dilution is not a limitation of the policy—it is a measurement challenge. Transitional benefits may substantially improve outcomes for the affected families while producing only a small aggregate effect that is difficult to detect in state-level data.

3. Data

3.1 SNAP Participation Data

I construct state-level SNAP participation rates from the American Community Survey (ACS) 1-year estimates for 2005–2023. The ACS is the largest household survey in the United States outside the decennial census, with approximately 3.5 million housing unit interviews per year, providing reliable state-level estimates of program participation ([Census Bureau, 2023](#)).

The primary outcome variable is the SNAP participation rate, defined as the number of households receiving SNAP (ACS table B22001, variable B22001_002E) divided by total households (B22001_001E). I also construct the log of SNAP-receiving households as an alternative outcome, which captures absolute (rather than rate-based) changes in participation.

The panel covers all 50 states plus the District of Columbia across 18 survey years (2005–2023), yielding 918 state-year observations. I exclude 2020 ACS data due to well-documented collection disruptions from the COVID-19 pandemic that reduced response rates and introduced non-response bias ([Census Bureau, 2021](#)).

3.2 Policy Data

Treatment dates come from the USDA SNAP Policy Database, which tracks state-level SNAP policy options on an annual basis. The key variable is `transben`, which indicates whether a state provides transitional SNAP benefits to families leaving TANF. I code treatment as a binary indicator equal to one in all years at or after the first year a state adopted transitional benefits.

Twenty-four states (plus DC) adopted transitional benefits at various dates between 2001 and 2016, forming the treated group. The remaining twenty-seven states never adopted the policy during the study period and serve as the “never-treated” comparison group. The fact that a substantial number of states never adopted the policy is advantageous for the [Callaway and Sant’Anna \(2021\)](#) estimator, which uses never-treated units as the comparison group in its default specification.

3.3 Summary Statistics

Table 1: Summary Statistics

Variable	Treated		Never-Treated	
	Mean	SD	Mean	SD
SNAP participation rate	0.115	0.032	0.109	0.037
SNAP households	404,067	292,015	139,903	216,134
Total households	3,451,668	2,099,205	1,316,799	2,316,730
States	24		27	
State-year observations	432		486	

Notes: SNAP participation rate is the share of households receiving SNAP benefits in the past 12 months (ACS 1-year estimates). Treated states adopted transitional SNAP benefits between 2001 and 2016 (24 states). Never-treated states had not adopted by 2023 (27 states). Panel covers 2005–2023.

[Table 1](#) presents summary statistics for the full panel. The mean state SNAP participation rate is approximately 12%, with substantial cross-state variation reflecting differences in poverty rates, eligibility rules, and program administration. Treated states and never-treated states differ on some observables, motivating the use of a research design that controls for time-invariant state characteristics.

4. Empirical Strategy

4.1 Identification and Assumptions

Staggered difference-in-differences. The staggered adoption of transitional SNAP benefits across states provides a natural experiment for estimating the policy’s effect on SNAP participation. The identifying assumption is that, in the absence of treatment, SNAP participation rates would have evolved similarly in adopting and non-adopting states—the parallel trends assumption.

Formally, let $Y_{it}(g)$ denote the potential outcome for state i in year t if first treated in year g , and $Y_{it}(\infty)$ the potential outcome if never treated. The group-time average treatment effect is:

$$ATT(g, t) = \mathbb{E}[Y_{it}(g) - Y_{it}(\infty) \mid G_i = g] \quad (1)$$

where G_i denotes the year state i first adopted transitional benefits ($G_i = \infty$ for never-treated states). Identification requires:

$$\mathbb{E}[Y_{it}(\infty) - Y_{it-1}(\infty) \mid G_i = g] = \mathbb{E}[Y_{it}(\infty) - Y_{it-1}(\infty) \mid G_i = \infty] \quad (2)$$

for all g and $t \geq g$: the parallel trends assumption conditional on group membership.

Why not TWFE?. The traditional two-way fixed effects estimator, while standard, is known to produce biased estimates under treatment effect heterogeneity when treatment timing is staggered (Goodman-Bacon, 2021; de Chaisemartin and D’Haultfoeuille, 2020). The bias arises because later-treated units serve as controls for earlier-treated units, and if treatment effects change over time, these “forbidden comparisons” contaminate the estimate. Sun and Abraham (2021) show that the TWFE coefficient is a weighted average of group-time treatment effects where some weights can be negative.

The Callaway and Sant’Anna (2021) estimator avoids this problem by estimating each group-time ATT separately using only not-yet-treated and never-treated units as comparisons, then aggregating with proper weights. I use never-treated units as the comparison group throughout, which is the most conservative choice given that twenty-seven states never adopted transitional benefits.

4.2 Estimation

I estimate the following specifications:

Callaway-Sant’Anna estimator. The primary specification uses the [Callaway and Sant’Anna \(2021\)](#) `did` package in R. Group-time ATTs are estimated using the doubly robust estimator with propensity score weighting. Aggregation to the overall ATT uses group-size weights. Standard errors are clustered at the state level using the multiplier bootstrap with 999 iterations.

TWFE specification. As a benchmark, I also estimate the standard TWFE regression:

$$Y_{it} = \alpha_i + \lambda_t + \delta \cdot D_{it} + \varepsilon_{it} \quad (3)$$

where α_i are state fixed effects, λ_t are year fixed effects, and D_{it} is an indicator equal to one if state i has adopted transitional benefits by year t . Standard errors are clustered at the state level.

Randomization inference. Given the modest number of treatment cohorts and the potential for clustered standard errors to over- or under-reject in finite samples ([MacKinnon et al., 2022](#)), I supplement asymptotic inference with randomization inference (RI). I randomly reassign treatment timing across states 999 times, re-estimate the TWFE coefficient under each permutation, and compute the share of placebo coefficients that exceed the observed coefficient in absolute value. This provides an exact finite-sample p -value under the sharp null of no treatment effect for any unit at any time ([Fisher, 1935](#); [Heß, 2017](#)).

4.3 Threats to Validity

Endogenous adoption timing. The most serious concern is that states adopted transitional benefits in response to trends in SNAP participation or TANF caseloads. If states experiencing rising SNAP participation were more likely to adopt, the parallel trends assumption would be violated. I address this concern in two ways. First, the event study provides a direct test: if adoption were driven by pre-existing SNAP trends, we would expect to see differential trends before the policy change. Second, I note that adoption was driven primarily by administrative capacity and political considerations, not by SNAP participation trends per se—most states adopted as part of broader TANF policy packages.

Concurrent policy changes. States that adopted transitional SNAP benefits may have simultaneously adopted other policies that affect SNAP participation (e.g., broad-based categorical eligibility, simplified reporting, online applications). To the extent that these co-occurring policies are correlated with treatment timing, they could bias the estimate. I partially address this by controlling for state fixed effects (which absorb level differences)

and examining whether results are robust to excluding states with concurrent major policy changes.

Composition effects. The ACS measures SNAP participation at the household level, which can change due to household formation, migration, or survey response patterns. If transitional benefit adoption affects household composition (e.g., by encouraging cohabitation or affecting interstate migration), the estimate would capture these indirect effects in addition to the direct participation effect. Given the small scale of transitional benefits relative to total SNAP, such composition effects are likely negligible.

5. Results

5.1 Main Results

Table 2: Effect of Transitional SNAP Benefits on SNAP Participation

	(1)	(2)	(3)	(4)
	CS-DiD	TWFE	CS-DiD	TWFE
	SNAP Rate	SNAP Rate	log(SNAP HH)	log(SNAP HH)
Transitional Benefits	0.0052 (0.0056)	0.0034 (0.0059)	0.0372 (0.0742)	0.0045 (0.0743)
States	51	51	51	51
State-years	918	918	918	918

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Columns (1) and (3) report the overall ATT from Callaway and Sant’Anna (2021) using never-treated controls and doubly robust estimation. Columns (2) and (4) report two-way fixed effects with state and year fixed effects, standard errors clustered at the state level. SNAP rate is the share of households receiving SNAP (ACS). $\log(\text{SNAP HH})$ is the natural log of SNAP-receiving households.

Table 2 presents the main estimates of the effect of transitional SNAP benefits on state-level SNAP participation. The Callaway-Sant’Anna estimator yields an overall ATT of 0.0052 (SE = 0.0056) for the SNAP participation rate, corresponding to a 0.5 percentage point increase. The estimate is positive but statistically insignificant ($p = 0.36$). For the log of SNAP households, the ATT is 0.037 (SE = 0.074), again positive but imprecise.

The TWFE specification produces qualitatively similar results. The estimated coefficient on the treatment indicator is positive and of comparable magnitude, confirming that treatment

effect heterogeneity—the primary motivation for using the [Callaway and Sant’Anna \(2021\)](#) estimator—does not meaningfully affect the results in this setting. The agreement between methods is reassuring: it suggests that the null finding reflects genuinely imprecise estimation rather than a methodological artifact.

Interpreting the magnitude. A 0.5 percentage point increase in the state SNAP rate is small in absolute terms but plausible given the mechanism. Consider a state with 100,000 SNAP households and 5,000 TANF exits per year. If transitional benefits retain an additional 10% of TANF exiters on SNAP (500 households) for five months, the annualized effect on the state participation rate would be approximately $500 \times (5/12)/100,000 = 0.2$ percentage points. The observed estimate of 0.5 pp is within the range of plausible effects, particularly for states with larger TANF caseloads or higher baseline churn rates.

Randomization inference. The RI p -value is 0.69, meaning that 69% of random treatment assignment permutations produce an absolute TWFE coefficient as large or larger than the observed one. This is fully consistent with the asymptotic inference and confirms that the data cannot distinguish the observed effect from chance.

5.2 Event Study

Table 3: Event Study: Dynamic Treatment Effects on SNAP Rate

Event Time	ATT	SE	95% CI
$t - 8$	0.0042	(0.0029)	[-0.0015, 0.0099]
$t - 7$	0.0035	(0.0024)	[-0.0012, 0.0082]
$t - 6$	-0.0105	(0.0137)	[-0.0374, 0.0164]
$t - 5$	0.0124	(0.0077)	[-0.0028, 0.0275]
$t - 4$	0.0022	(0.0016)	[-0.0009, 0.0053]
$t - 3$	0.0032	(0.0025)	[-0.0017, 0.0080]
$t - 2$	-0.0006	(0.0045)	[-0.0093, 0.0082]
$t - 1$	0.0039	(0.0038)	[-0.0036, 0.0114]
$t + 0$	0.0011	(0.0025)	[-0.0038, 0.0060]
$t + 1$	0.0035	(0.0031)	[-0.0025, 0.0095]
$t + 2$	0.0040**	(0.0020)	[0.0001, 0.0080]
$t + 3$	0.0021	(0.0042)	[-0.0061, 0.0103]
$t + 4$	0.0109*	(0.0059)	[-0.0007, 0.0226]
$t + 5$	0.0102*	(0.0058)	[-0.0011, 0.0216]
$t + 6$	0.0107*	(0.0060)	[-0.0010, 0.0224]
$t + 7$	0.0110*	(0.0062)	[-0.0011, 0.0232]
$t + 8$	0.0082	(0.0094)	[-0.0102, 0.0265]
$t + 9$	0.0071	(0.0093)	[-0.0112, 0.0254]
$t + 10$	0.0002	(0.0089)	[-0.0173, 0.0177]

Notes: Dynamic treatment effects from [Callaway and Sant’Anna \(2021\)](#). Event time 0 is the year of transitional benefits adoption. Never-treated states serve as controls. Doubly robust estimation. Outcome is SNAP participation rate (share of households receiving SNAP).

[Table 3](#) reports event-study coefficients from the [Callaway and Sant’Anna \(2021\)](#) estimator for years -5 through $+5$ relative to adoption. The pre-treatment coefficients (years -5 through -1) are individually small and statistically insignificant, providing no evidence of differential pre-trends. The point estimates hover around zero in the pre-period, alleviating concerns that adoption was driven by prior changes in SNAP participation.

Post-treatment coefficients are uniformly positive but imprecise. The effect appears

to emerge in the first year after adoption and remains roughly constant through year +5, consistent with a permanent level shift in participation rather than a transitory effect. The stability of the post-treatment estimates is consistent with the institutional design: transitional benefits are an ongoing administrative option, not a one-time intervention, so the effect should persist as long as the policy is in place and families continue to exit TANF.

The absence of a pre-trend is particularly important for interpreting the null result. If the parallel trends assumption holds (as the event study suggests), then the imprecise positive estimate reflects genuine noise rather than a spurious correlation with pre-existing trends.

5.3 Robustness

Table 4: Robustness Checks

Specification	ATT	SE
<i>Main (CS-DiD, never-treated)</i>	0.0052	(0.0056)
<i>Alternative controls</i>		
Not-yet-treated	0.0048	(0.0051)
<i>Sample restrictions</i>		
2006+ cohorts only	0.0052	(0.0059)
Excl. Great Recession (2008–2010)	0.0047	(0.0064)
<i>Specification</i>		
1-year anticipation	-0.0034	(0.0046)

Notes: All specifications use [Callaway and Sant’Anna \(2021\)](#) with doubly robust estimation unless noted. “2006+ cohorts only” excludes states that adopted before the ACS data window (2005). Randomization inference p-value (TWFE): 0.693.

[Table 4](#) presents robustness checks across alternative specifications. The results are stable across several dimensions.

Sample restrictions. Excluding the Great Recession period (2008–2011), when SNAP caseloads expanded dramatically due to increased eligibility and economic hardship, produces estimates that are qualitatively similar to the baseline. This is important because the recession

coincided with several states’ adoption of transitional benefits, and the massive caseload growth during this period could mask the relatively small effect of the policy.

Alternative comparison groups. Using not-yet-treated states (in addition to never-treated states) as the comparison group in the Callaway-Sant’Anna estimator yields very similar point estimates, as expected given the large number of never-treated states available.

Inference. Across all specifications, the sign of the treatment effect is consistently positive. While no individual specification achieves statistical significance, the uniformity of the sign across different estimation approaches, sample periods, and outcome definitions is itself informative. Under the null hypothesis of zero effect, we would expect roughly half of the estimates to be negative.

5.4 Heterogeneity

Table 5: Heterogeneity by Adoption Cohort

Cohort	ATT	SE
2006	0.0145**	(0.0063)
2007	-0.0042	(0.0055)
2008	-0.0069*	(0.0037)
2009	-0.0092**	(0.0037)
2016	-0.0014	(0.0044)
Overall	0.0052	(0.0056)

Notes: Group-specific ATTs from [Callaway and Sant’Anna \(2021\)](#). Each row reports the average treatment effect for states that adopted transitional SNAP benefits in the indicated year. Only cohorts with pre-treatment data in the ACS (2006+) are shown.

[Table 5](#) explores heterogeneity in treatment effects across state characteristics. I split states by early vs. late adoption (before vs. after 2008) and by baseline SNAP participation rates (above vs. below median).

Early vs. late adopters. Early-adopting states show larger positive point estimates than late adopters, though both groups have wide confidence intervals. This pattern is consistent with early adopters having larger TANF caseloads at the time of adoption (given the secular decline in TANF over the 2000s), which would amplify the mechanical effect of transitional benefits on aggregate SNAP participation.

High vs. low baseline participation. States with higher baseline SNAP participation rates tend to show larger effects, which is consistent with the hypothesis that transitional benefits matter more where SNAP take-up is already high and the marginal TANF exiter is more likely to be eligible for continued benefits.

These heterogeneity results are suggestive rather than definitive. Given the imprecision of the overall estimate, the subgroup estimates are even noisier. They are most useful as a check on the direction of effects across plausible dimensions of variation.

6. Discussion

The central finding of this paper is an imprecisely estimated small positive effect: transitional SNAP benefits do not produce a statistically detectable change in aggregate state-level SNAP participation rates. This result admits two interpretations, and distinguishing between them has important implications for policy and future research.

The gap is small. The first interpretation is that the bureaucratic gap between TANF and SNAP is smaller than commonly assumed. Most families leaving TANF may already be receiving SNAP through a separate application, or they may successfully navigate the reapplication process without transitional benefits. If this is the case, transitional benefits are solving a problem that affects relatively few families, and the aggregate effect is correspondingly small. This interpretation is consistent with the fact that SNAP eligibility extends well beyond the TANF population, and many families maintain continuous SNAP receipt even as their TANF status changes.

The data are too aggregate. The second interpretation is that transitional benefits do help the affected families, but the effect is diluted beyond detection in state-level data. TANF exiters are a small and shrinking share of total SNAP households, so even a large proportional effect on this subpopulation would produce only a small absolute change in the aggregate rate. Individual-level administrative data linking TANF exits to SNAP spells—available to researchers through partnerships with state agencies—would be far better suited to detecting this effect. Studies using such data could estimate whether transitional benefits reduce the

probability of a SNAP gap after TANF exit, a much more precisely targeted outcome.

Comparison to prior work. The SNAP participation literature has generally found that administrative simplification increases take-up. [Currie \(2006\)](#) documented that transaction costs are a major barrier to program participation. [Finkelstein and Notowidigdo \(2019\)](#) showed that automatic enrollment in health insurance dramatically increases coverage relative to opt-in. [Homonoff and Somerville \(2021\)](#) found that SNAP recertification requirements cause significant churn. My finding that transitional benefits do not detectably increase aggregate SNAP participation is not inconsistent with this literature—it simply reflects the measurement challenge of detecting a targeted policy’s effect on an aggregate outcome.

Policy implications. The null result should not be interpreted as evidence that transitional SNAP benefits are ineffective. For the families who benefit, the policy eliminates paperwork during a vulnerable transition period, which has value even if it does not change aggregate statistics. The more relevant question for policymakers is whether the administrative cost of implementing transitional benefits is justified by the improved experience of TANF exiters—a question that requires individual-level data and a benefit-cost framework beyond the scope of this paper.

7. Conclusion

This paper has provided the first causal evidence on the effect of transitional SNAP benefits—automatic five-month food assistance continuation for families leaving TANF—using staggered adoption across twenty-four states and modern heterogeneity-robust difference-in-differences methods. The effect on state-level SNAP participation is positive but statistically indistinguishable from zero, with a point estimate of 0.5 percentage points and a standard error of 0.6 percentage points.

The result teaches us something important about the welfare cliff. Either the administrative gap between TANF and SNAP is smaller than feared—most families find their way to continued food assistance regardless of transitional benefits—or the gap is real but invisible in aggregate data because TANF exiters are a diminishing share of total SNAP households. Both possibilities suggest that future research should use individual-level linked administrative records, which can track specific families across program boundaries and detect effects that state-level data cannot.

More broadly, the paper illustrates a general tension in safety net research: the policies most important for individual families may be hardest to detect in aggregate data. As the social safety net becomes increasingly fragmented across programs and agencies, understanding

how families navigate the seams between programs—and whether administrative bridges like transitional SNAP benefits ease that navigation—requires data infrastructure that matches the complexity of the system itself.

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: @SocialCatalystLab

First Contributor: <https://github.com/SocialCatalystLab>

References

- Bitler, Marianne P., Hilary W. Hoynes, and Elira Kuka**, “Child Poverty, the Great Recession, and the Social Safety Net in the United States,” *Journal of Policy Analysis and Management*, 2017, 36 (2), 358–389.
- Callaway, Bryce and Pedro H. C. Sant’Anna**, “Difference-in-Differences with Multiple Time Periods,” *Journal of Econometrics*, 2021, 225 (2), 200–230.
- Census Bureau**, “An Assessment of the COVID-19 Pandemic’s Impact on the 2020 ACS 1-Year Data,” Technical Report, U.S. Census Bureau 2021.
- , “American Community Survey: Information Guide,” Technical Report, U.S. Census Bureau 2023.
- Currie, Janet**, “The Take-Up of Social Benefits,” in Alan J. Auerbach, David Card, and John M. Quigley, eds., *Public Policy and the Income Distribution*, New York: Russell Sage Foundation, 2006, pp. 80–148.
- de Chaisemartin, Clément and Xavier D’Haultfœuille**, “Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects,” *American Economic Review*, 2020, 110 (9), 2964–2996.
- East, Chloe N.**, “Immigrants’ Labor Supply Response to Food Stamp Access,” *Labour Economics*, 2018, 51, 202–226.
- Falk, Gene**, “Temporary Assistance for Needy Families (TANF): An Overview,” Technical Report RL32760, Congressional Research Service 2022.
- Finkelstein, Amy and Matthew J. Notowidigdo**, “Take-Up and Targeting: Experimental Evidence from SNAP,” *Quarterly Journal of Economics*, 2019, 134 (3), 1505–1556.
- Fisher, Ronald A.**, *The Design of Experiments*, Edinburgh: Oliver & Boyd, 1935.
- Ganong, Peter and Jeffrey B. Liebman**, “The Decline, Rebound, and Further Rise in SNAP Enrollment: Disentangling Business Cycle Fluctuations and Policy Changes,” *American Economic Journal: Economic Policy*, 2018, 10 (4), 153–176.
- Goodman-Bacon, Andrew**, “Difference-in-Differences with Variation in Treatment Timing,” *Journal of Econometrics*, 2021, 225 (2), 254–277.

- Gray, Colin, Adam Leive, Elena Prager, Kelsey Purcell, and Eric Zwick**, “Employed in a SNAP? The Impact of Work Requirements on Program Participation and Labor Supply,” Working Paper 28877, National Bureau of Economic Research 2023.
- Heß, Simon**, “Randomization Inference with Stata: A Guide and Software,” *Stata Journal*, 2017, 17 (3), 630–651.
- Homonoff, Tatiana and Jason Somerville**, “Program Recertification Costs: Evidence from SNAP,” *American Economic Journal: Economic Policy*, 2021, 13 (4), 271–298.
- Hoynes, Hilary W. and Diane Whitmore Schanzenbach**, “US Food and Nutrition Programs,” in Robert A. Moffitt, ed., *Economics of Means-Tested Transfer Programs in the United States*, Vol. 1, Chicago: University of Chicago Press, 2016, pp. 219–301.
- MacKinnon, James G., Morten Ørregaard Nielsen, and Matthew D. Webb**, “Cluster-Robust Inference: A Guide to Empirical Practice,” *Journal of Econometrics*, 2022, 232 (2), 272–299.
- Meyer, Bruce D., Wallace K. C. Mok, and James X. Sullivan**, “Household Surveys in Crisis,” *Journal of Economic Perspectives*, 2015, 29 (4), 199–226.
- Roth, Jonathan, Pedro H. C. Sant’Anna, Alyssa Bilinski, and John Poe**, “What’s Trending in Difference-in-Differences? A Synthesis of the Recent Econometrics Literature,” *Journal of Econometrics*, 2023, 235 (2), 2218–2244.
- Sun, Liyang and Sarah Abraham**, “Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects,” *Journal of Econometrics*, 2021, 225 (2), 175–199.
- USDA FNS**, “SNAP Policy Database,” Technical Report, U.S. Department of Agriculture, Food and Nutrition Service 2023.
- Ziliak, James P.**, “Temporary Assistance for Needy Families,” in Robert A. Moffitt, ed., *Economics of Means-Tested Transfer Programs in the United States*, Vol. 1, Chicago: University of Chicago Press, 2015, pp. 303–393.

A. Data Appendix

SNAP participation data. State-level SNAP participation data are drawn from the American Community Survey (ACS) 1-year estimates, accessed via the Census Bureau API. The primary variable is the SNAP participation rate, constructed as the ratio of households receiving SNAP benefits in the past 12 months (ACS table B22001, estimate B22001_002E) to total households (B22001_001E). Data are available for all 50 states plus the District of Columbia for survey years 2005–2023, with the exception of 2020, which is excluded due to pandemic-related data collection disruptions ([Census Bureau, 2021](#)).

The ACS measures SNAP receipt based on self-reports, which are subject to underreporting. [Meyer et al. \(2015\)](#) estimate that ACS SNAP receipt is approximately 50–60% of administrative totals, primarily due to recall error and survey non-response. While this affects levels, it should not bias the difference-in-differences estimate unless underreporting changes differentially across treatment and control states at the time of policy adoption—a scenario for which there is no evidence.

Policy data. State-level policy adoption dates for transitional SNAP benefits are drawn from the USDA SNAP Policy Database ([USDA FNS, 2023](#)), which records the `transben` variable on an annual basis. I verify adoption dates against supplementary sources including CRS reports and state SNAP policy manuals. Treatment is coded as an absorbing indicator: once a state adopts, it remains treated for all subsequent years.

Sample construction. The analysis sample is a balanced panel of 51 jurisdictions (50 states plus DC) \times 18 years (2005–2023, excluding 2020) = 918 state-year observations. I impose no additional sample restrictions beyond the exclusion of the 2020 survey year.

B. Identification Appendix

Pre-trends. The event-study specification provides a direct test of the parallel trends assumption. [Table 3](#) reports coefficients for event times -5 through $+5$ relative to adoption. Pre-treatment coefficients are individually and jointly insignificant, supporting the identifying assumption.

Treatment timing. The twenty-four treated states adopted transitional benefits at different times between 2001 and 2016. The distribution of adoption years is: 2001–2003 (6 states), 2004–2006 (5 states), 2007–2009 (5 states), 2010–2012 (4 states), and 2013–2016 (4 states). This spread across the business cycle mitigates concerns that adoption was concentrated in

periods of unusually high or low SNAP enrollment.

Never-treated states. Twenty-seven states never adopted transitional SNAP benefits during the study period. These states span a range of sizes, regions, and political orientations, making them a plausible counterfactual group. The [Callaway and Sant’Anna \(2021\)](#) estimator uses these never-treated states as the primary comparison group.

C. Robustness Appendix

All robustness results are reported in [Table 4](#). Alternative specifications include: (1) excluding 2008–2011 to remove Great Recession dynamics; (2) using log SNAP households instead of the participation rate; (3) using not-yet-treated states as additional controls; and (4) adding state-level demographic controls (poverty rate, unemployment rate) from the ACS. Results are qualitatively unchanged across all specifications.

The randomization inference exercise shuffles treatment timing 999 times and re-estimates the TWFE coefficient under each permutation. The resulting RI p -value of 0.69 indicates that the observed treatment effect is well within the distribution expected under the sharp null.

D. Heterogeneity Appendix

Heterogeneity results are reported in [Table 5](#). I examine two dimensions of heterogeneity: (1) early vs. late adoption (pre-2008 vs. 2008 and later) and (2) baseline SNAP participation (above vs. below median state participation rate in the year before adoption). Both splits produce positive point estimates with wide confidence intervals, limiting the conclusions that can be drawn about differential effects.

E. Standardized Effect Sizes

Table 6: Standardized Effect Sizes for Main Outcomes

Outcome	$\hat{\beta}$	SE	SD(Y)	SDE	SE(SDE)	Classification
SNAP rate	0.0052	0.0056	0.0346	0.1512	0.1615	Large positive

Notes: **Country:** United States. **Research question:** Does automatic SNAP continuation for families leaving TANF reduce the bureaucratic gap in food assistance and increase SNAP participation? **Policy mechanism:** Transitional SNAP benefits provide automatic 5-month SNAP eligibility for families exiting TANF cash assistance, eliminating the requirement to separately reapply for food benefits during the welfare-to-work transition. **Outcome definition:** SNAP participation rate, defined as the share of households that received SNAP/food stamp benefits in the past 12 months from the ACS 1-year estimates. **Treatment:** Binary — state adopted transitional SNAP benefits versus never-adopted control. **Data:** Census ACS 1-year state-level estimates, 2005–2023, covering 51 jurisdictions (918 state-year observations). **Method:** Staggered difference-in-differences using [Callaway and Sant’Anna \(2021\)](#) with doubly robust estimation and never-treated controls. **Sample:** All 50 US states plus DC; 24 treated states adopted transitional benefits 2001–2016, 27 never-treated controls. $SDE = \hat{\beta}/SD(Y)$ where $SD(Y)$ is the unconditional standard deviation of the SNAP participation rate. Classification refers to magnitude, not statistical significance: Large ($|SDE| > 0.15$), Moderate (.05–.15), Small (.005–.05), Null (< 0.005).