

Back to Work? Early Termination of Pandemic Unemployment Benefits and Medicaid Home Care Provider Supply*

APEP Autonomous Research[†]

@ai1scl

@SocialCatalystLab

March 2, 2026

Abstract

The \$300/week Federal Pandemic Unemployment Compensation supplement often exceeded wages for Medicaid home care aides, potentially reducing provider supply during a period of acute need. I exploit the staggered early termination of these benefits across 26 U.S. states in June–July 2021 to estimate the effect on home and community-based services (HCBS) provider supply using novel T-MSIS claims data covering 227 million Medicaid billing records. Callaway-Sant’Anna difference-in-differences estimates show that early termination increased active HCBS providers by 6.3 percent and beneficiaries served by 14.9 percent. A placebo test on higher-wage behavioral health providers yields a precise null, consistent with a labor supply mechanism operating through the reservation wage. Randomization inference, within-region comparisons, and a triple-difference design corroborate these findings.

JEL Codes: I13, J22, J38, H75

Keywords: unemployment insurance, Medicaid, HCBS, provider supply, labor supply, COVID-19

*This paper is a revision of APEP-0448. See https://github.com/SocialCatalystLab/ape-papers/tree/main/apep_0448/v1/paper.pdf for the original.

[†]Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch

1. Introduction

When the federal government offered unemployed Americans an extra \$300 per week during the pandemic, it created an awkward arithmetic for home care workers. A personal care aide earning \$12 an hour—the median wage for the occupation (PHI, 2021)—could collect more by staying home than by assisting a disabled Medicaid beneficiary with bathing, dressing, or medication management. The resulting tension between income support and care delivery became one of the most politically charged labor market debates of 2021: were enhanced unemployment benefits keeping workers on the sidelines while vulnerable people went without essential services?

Twenty-six state governors answered affirmatively. Between June 12 and July 31, 2021, these states—overwhelmingly Republican-led—voluntarily terminated the Federal Pandemic Unemployment Compensation (FPUC) supplement before its scheduled federal expiration on September 6. Their stated rationale was that generous benefits were discouraging workers from returning to jobs, particularly in low-wage sectors where the supplement exceeded market wages (Holzer et al., 2021). The remaining 24 states and the District of Columbia maintained benefits through the federal deadline.

This paper asks whether the early termination of pandemic unemployment benefits restored Medicaid home and community-based services (HCBS) provider supply. HCBS—personal care, habilitation, and attendant care services delivered in beneficiaries’ homes—constitutes the backbone of long-term care for elderly and disabled Medicaid enrollees. The sector employs roughly 4.6 million direct care workers nationally, the vast majority of whom earn wages in the \$12–\$15 per hour range (PHI, 2021). This wage structure made HCBS uniquely vulnerable to the UI work disincentive: the \$300 weekly supplement alone equaled \$7.50 per hour for a full-time worker, and combined with state UI benefits of \$200–\$400 per week, total UI income routinely exceeded HCBS wages.

I exploit the staggered timing of early termination across 26 states to estimate its causal effect on HCBS provider supply using a novel administrative dataset. The Transformed Medicaid Statistical Information System (T-MSIS) Provider Spending file, released by the Department of Health and Human Services in February 2026, contains 227 million billing records covering every Medicaid claim from January 2018 through December 2024. By linking billing NPIs to the National Plan and Provider Enumeration System (NPPES), I construct a state-by-month panel of active HCBS providers, total claims, Medicaid payments, and unique beneficiaries served.

The research design is a staggered difference-in-differences using the Callaway and Sant’Anna (2021) heterogeneity-robust estimator. Treatment cohorts are defined by the first

full month of exposure: 22 states enter treatment in July 2021, and 4 states enter in August 2021. The 25 never-treated jurisdictions serve as the comparison group. With 41 months of pre-treatment data (January 2018 through May 2021), the design offers exceptional power for parallel-trend testing.

The results reveal a substantial and statistically significant effect. The Callaway-Sant’Anna simple aggregated ATT shows that early UI termination increased the number of active HCBS providers by 6.3 percent (SE = 0.029). The effect on beneficiaries served is even larger: a 14.9 percent increase (SE = 0.068), suggesting that returning providers served multiple beneficiaries. Total claims increased by 8.7 percent, though this estimate is imprecise. The event study shows negligible pre-trends and a gradual post-treatment increase that stabilizes approximately six months after termination, consistent with a labor supply response that takes time to translate into Medicaid billing.

A placebo test provides critical support for the labor supply mechanism. Behavioral health providers—who bill H-codes for services like community psychiatric rehabilitation and crisis intervention—earn substantially higher wages (\$18–\$25 per hour) than HCBS aides. If the effect operates through the reservation wage channel, UI termination should have little impact on behavioral health providers, whose wages already exceeded UI income. The data confirm this prediction: the CS-DiD ATT for behavioral health providers is 0.8 percent (SE = 0.051), a precise null that contrasts sharply with the 6.3 percent HCBS effect.

The findings survive an extensive battery of robustness checks. A Bacon decomposition reveals that 99.4 percent of the two-way fixed effects estimate derives from comparisons between treated and untreated states, with negligible contamination from timing variation. Within-region analysis restricted to Southern states—where both treated and untreated states are geographically proximate—yields a nearly identical estimate ($\beta = 0.111$, SE = 0.055). A placebo test shifting treatment timing two years earlier (to 2019) produces a null effect ($\beta = 0.041$, $p = 0.38$). Randomization inference, permuting treatment assignment and re-estimating the CS-DiD ATT for each permutation, yields a p -value of 0.040, rejecting the sharp null at the 5% level.

This paper contributes to three literatures. First, it extends the growing body of evidence on the employment effects of pandemic UI generosity (Holzer et al., 2021; Coombs et al., 2022; Ganong et al., 2022; Dube, 2021). These studies measure aggregate employment responses using Current Population Survey or unemployment insurance claims data. I show that the effects extend to a specific, policy-critical sector—Medicaid home care—where the welfare implications differ fundamentally from aggregate labor market statistics. A worker returning to a warehouse job is a labor market outcome; a worker returning to assist a disabled person with daily living activities is a healthcare access outcome.

Second, the paper contributes to the nascent literature on Medicaid provider supply dynamics (Dague and Ukert, 2023; Zuckerman et al., 2021). Most Medicaid provider research focuses on physician participation in fee-for-service programs. The HCBS workforce—which serves the most vulnerable Medicaid beneficiaries—has been largely invisible to researchers because provider-level billing data was unavailable until the T-MSIS release. This paper demonstrates that HCBS provider supply responds to labor market incentives, a finding with direct implications for workforce policy.

Third, the paper provides the first evidence on the sector-specific effects of UI generosity in healthcare, connecting the unemployment insurance literature to the healthcare labor economics literature (Staiger et al., 2010; Autor et al., 2020). The finding that enhanced UI benefits can reduce healthcare workforce participation among the lowest-paid providers has implications for the design of future income support programs, particularly in sectors where labor supply directly affects vulnerable populations’ access to essential services.

2. Institutional Background

2.1 Federal Pandemic Unemployment Benefits

The Coronavirus Aid, Relief, and Economic Security (CARES) Act, signed March 27, 2020, created the Federal Pandemic Unemployment Compensation (FPUC) program, which provided a flat \$600 weekly supplement to all UI recipients regardless of their pre-separation earnings. When the initial authorization expired in July 2020, political negotiations produced a gap followed by a reduced \$300 weekly supplement under the Continued Assistance Act (December 2020) and the American Rescue Plan Act (March 2021), which extended benefits through September 6, 2021.

The \$300 weekly supplement was not trivial relative to low-wage workers’ earnings. For a full-time worker earning \$12 per hour, the supplement alone equaled 62.5 percent of weekly wages. Combined with state UI benefits—which averaged approximately \$300 per week nationally but varied from \$190 in Mississippi to \$550 in Massachusetts—total UI income could reach \$600–\$850 per week, equivalent to \$15–\$21 per hour at 40 hours. For HCBS workers in low-benefit states, UI income roughly matched or modestly exceeded their employment income. In higher-benefit states, UI income substantially exceeded HCBS wages.

Beginning in May 2021, governors of 26 states announced their intention to terminate one or more federal pandemic UI programs early. The first terminations took effect June 12 (Alaska, Iowa, Mississippi, Missouri), with subsequent waves on June 19 (8 states), June 26 (8 states), and scattered dates through July 31 (Louisiana was last). Table 5 provides the complete timeline. Twenty-five of the 26 states had Republican governors; Louisiana’s

Democratic governor was the sole exception. The primary stated justification was that enhanced benefits were impeding labor market recovery, particularly in hospitality, retail, and personal services ([Holzer et al., 2021](#)).

Federal courts briefly intervened in several states. Indiana’s termination was temporarily blocked by a state court injunction in June 2021, though the ruling was later reversed and the termination took effect June 19. Maryland terminated FPUC on July 3 but reinstated PUA and PEUC (extended duration programs) later that month; crucially, the \$300 weekly FPUC supplement—the binding constraint for HCBS workers—was not reinstated. For the main analysis, I code Maryland as treated effective August 2021 based on its FPUC termination date. Robustness checks excluding Maryland yield virtually identical results (not tabulated).

2.2 The HCBS Workforce

Home and community-based services (HCBS) represent the largest category of Medicaid long-term services and supports. These services—personal care assistance, habilitation, attendant care, and respite care—enable elderly and disabled individuals to live in their homes and communities rather than in institutional settings. In 2020, Medicaid spending on HCBS exceeded \$100 billion, surpassing institutional spending for the first time ([MACPAC, 2022](#)).

The HCBS workforce consists primarily of direct care workers: personal care aides, home health aides, and direct support professionals. The Bureau of Labor Statistics reports approximately 4.6 million workers in these occupations nationally. Their compensation profile is striking in its uniformity and low level. The median hourly wage for home health and personal care aides was \$14.15 in 2021, with the 25th percentile at \$11.37 ([Bureau of Labor Statistics, 2022](#)). Benefits are uncommon: fewer than 40 percent receive employer-sponsored health insurance, and paid leave is rare ([PHI, 2021](#)).

The workforce was already in crisis before the pandemic. Annual turnover rates for direct care workers range from 40 to 80 percent across studies, driven by low wages, physically demanding work, irregular schedules, and limited career advancement opportunities ([Scales, 2020](#)). The pandemic intensified these challenges: workers faced infection risk from close physical contact with vulnerable clients, childcare disruptions, and—crucially—the availability of more generous income through unemployment insurance.

Unlike behavioral health services, which received emergency telehealth waivers during the pandemic, HCBS requires physical presence. A personal care aide cannot assist with bathing, dressing, or meal preparation remotely. This structural feature of the work creates a sharp distinction: when UI benefits exceeded HCBS wages, workers had to choose between in-person service delivery and staying home with higher income. The early termination of

UI benefits altered this calculus by reducing the outside option, potentially pushing workers back into the HCBS labor market.

2.3 Political Economy of Early Termination

The decision to terminate early was neither random nor primarily motivated by healthcare concerns. Twenty-five of the 26 early-terminating states had Republican governors; Louisiana’s Democratic governor was the exception. The timing clustered tightly: 22 states terminated effective June 12–30, with only four states (Maryland, Tennessee, Arizona, and Louisiana) acting in July 2021. This clustering creates the two treatment cohorts used in the empirical design.

Several features of the political economy are relevant for identification. First, governors cited general labor shortages—primarily in hospitality, food service, and construction—not healthcare workforce concerns. This is important because it implies that the “treatment” was not a response to HCBS-specific trends, reducing concerns about reverse causality. Second, the partisan alignment suggests that termination decisions were driven by ideology rather than local economic conditions, which supports the plausibility of parallel trends conditional on time and state fixed effects. Third, the near-simultaneous timing among most early terminators means that cross-state variation in treatment timing is limited, which motivates the Callaway-Sant’Anna approach with its explicit cohort structure over a simple TWFE that would exploit within-treated timing variation.

The early termination also encompassed programs beyond FPUC. Most early-terminating states simultaneously ended Pandemic Unemployment Assistance (PUA), which covered gig workers and self-employed individuals, and Pandemic Emergency Unemployment Compensation (PEUC), which extended the duration of regular benefits. While my analysis focuses on the FPUC supplement—the most relevant margin for HCBS workers, who typically qualified for regular UI—the simultaneous withdrawal of all pandemic programs may have amplified the effect by eliminating fallback options for workers who exhausted regular benefits.

3. Conceptual Framework

Consider a direct care worker choosing between employment in HCBS and non-employment (collecting UI benefits). The worker accepts employment if the market wage w exceeds the reservation wage w^* , which depends on non-labor income including UI benefits:

$$\text{Work if } w \geq w^* = f(b_{UI}, b_{state}, \theta) \tag{1}$$

where b_{UI} is the federal UI supplement (\$300/week under FPUC), b_{state} is the state UI benefit, and θ captures individual characteristics (disutility of work, household composition, health risk tolerance).

When $b_{UI} > 0$, the reservation wage rises. If $w^* > w$ for a non-trivial share of potential HCBS workers, some exit the labor force or delay return to work. The early termination of FPUC sets $b_{UI} = 0$ in treated states, reducing w^* and inducing a flow of workers back into HCBS employment.

Three predictions follow:

Prediction 1: Extensive margin response. Early UI termination increases the number of active HCBS providers (NPIs billing T-codes to Medicaid), as workers whose reservation wage previously exceeded the HCBS wage return to billing.

Prediction 2: Effect concentrated in low-wage services. The effect should be largest for personal care and attendant care (T1019, S5125)—the lowest-paid HCBS services—and smallest for services where provider wages exceed the UI threshold. Behavioral health providers (H-codes), who earn \$18–\$25 per hour, should show negligible effects because their wages already exceeded UI income even with the supplement.

Prediction 3: Gradual onset. The provider supply response should appear gradually rather than instantaneously, because (a) returning to Medicaid billing requires reactivation of provider agreements, (b) clients may need reassignment, and (c) administrative lags in claims processing create measurement delays.

4. Data

4.1 T-MSIS Medicaid Provider Spending

The primary data source is the Transformed Medicaid Statistical Information System (T-MSIS) Provider Spending file, released by the Department of Health and Human Services on February 9, 2026 ([Department of Health and Human Services, 2026](#)). This dataset represents the first public release of provider-level Medicaid billing data and contains 227 million records covering all 50 states, the District of Columbia, and U.S. territories from January 2018 through December 2024.

Each observation is uniquely identified by the combination of billing NPI, servicing NPI, HCPCS procedure code, and month of service. The key outcome variables are the number of unique beneficiaries served, total claims submitted, and total Medicaid payments. Critically, the time variable (`CLAIM_FROM_MONTH`) records the date of service, not the date of claim submission or adjudication, which ensures that outcomes are correctly aligned with the policy timeline.

I define HCBS providers as those billing codes beginning with “T” (personal care, habilitation, home health) and selected “S” codes (S5125 for attendant care, S5130, S5150 for respite care). These Medicaid-specific codes have no Medicare equivalent and capture the core HCBS workforce. For the placebo analysis, I define behavioral health providers as those billing “H” codes (community psychiatric rehabilitation, crisis intervention, behavioral health day treatment).

The dataset has important limitations. It contains no state identifier, provider specialty, or beneficiary demographics. Geographic assignment requires linking billing NPIs to the National Plan and Provider Enumeration System (NPPES), which provides practice state and ZIP code. The NPPES match rate for billing NPIs is 98.1 percent in our sample, indicating minimal loss from geographic linkage.

Cell suppression removes rows with fewer than 12 total claims, which disproportionately affects small-volume providers in rural areas. This censoring is unlikely to bias the main results, as the analysis operates at the state-month level where aggregation smooths individual-provider suppression.

Like most administrative claims datasets, T-MSIS exhibits a reporting lag: the most recent months contain incomplete data as claims continue to be submitted and adjudicated. In our data, December 2024 shows a sharp decline in provider counts (mean of 234 vs. 344 in October), clearly reflecting incomplete reporting rather than a real workforce contraction. All figures and descriptive analyses truncate the sample at November 2024 to avoid this artifact. The regression estimates are unaffected because the treatment window (June–August 2021) is well within the period of complete data.

4.2 NPPES and Geographic Assignment

The NPPES bulk extract (January 2026 vintage) contains 9.3 million registered NPIs with practice location, provider taxonomy, entity type, and lifecycle dates. I link each T-MSIS billing NPI to NPPES using an exact NPI match, obtaining practice state for 98.1 percent of HCBS billing NPI-month observations.

4.3 Treatment Data

Early UI termination dates come from Ballotpedia’s state-by-state tracker ([Ballotpedia, 2021](#)). I code each state’s termination date and define the first full month of exposure as the calendar month following the termination date (since most terminations occurred mid-month). This yields two treatment cohorts: 22 states entering treatment in July 2021 and 4 states entering in August 2021. The 25 remaining jurisdictions (24 states plus DC) serve as never-treated.

4.4 Panel Construction

I aggregate T-MSIS billing records to the state-by-month level separately for HCBS (T/S-codes) and behavioral health (H-codes). For each state-month cell, I compute four outcomes: the count of unique billing NPIs (active providers), total claims, total Medicaid payments, and unique beneficiaries served. The resulting panel covers 51 jurisdictions over 84 months (January 2018 through December 2024), yielding 4,284 state-month observations for HCBS and 4,284 for behavioral health.

Several sample restrictions warrant discussion. First, I use the billing NPI (rather than the servicing NPI) as the provider identifier because it determines the entity responsible for Medicaid reimbursement and therefore most directly reflects provider participation decisions. In many states, a home care agency holds the billing NPI while individual aides hold servicing NPIs. The billing NPI thus captures both independent providers and agencies, providing a more complete measure of active care delivery organizations. Second, I exclude U.S. territories (Puerto Rico, Guam, Virgin Islands, American Samoa, Northern Mariana Islands) because these jurisdictions have fundamentally different Medicaid financing structures and were not subject to the same FPUC framework. Third, all outcomes are measured in logs to accommodate the large variation in program size across states—New York’s HCBS program is orders of magnitude larger than Wyoming’s—and to interpret coefficients as proportional effects.

The panel is balanced: every jurisdiction-month cell contains non-missing values for all four outcomes. Zero-provider months are rare (occurring primarily in small states during the earliest months of the sample) and are replaced with a value of 1 before taking logs to avoid undefined values. This affects fewer than 0.5 percent of observations.

4.5 Summary Statistics

Table 1 reports pre-treatment means for early-terminating and non-terminating states. The average early-terminating state has 318 active HCBS providers per month with \$29.5 million in Medicaid payments, compared to 287 providers and \$69.5 million in non-terminating states. The higher per-state spending in non-terminating states reflects the inclusion of large Medicaid programs (New York, California, Illinois) in that group. Behavioral health provider counts are smaller in magnitude but follow a similar pattern.

Table 1: Summary Statistics: Pre-Treatment Period (January 2018 – May 2021)

	Early Terminators		Maintained Benefits	
	Mean	SD	Mean	SD
<i>Panel A: HCBS (T/S-codes)</i>				
Active providers	318	402	287	263
Total claims	318,293	625,225	493,359	1,044,807
Total paid (\$M)	29.5	38.5	69.5	153.9
Unique beneficiaries	47,616	52,340	98,144	215,280
<i>Panel B: Behavioral Health (H-codes)</i>				
Active providers	199	238	332	291
N states	26		25	
N months	41		41	

Notes: Pre-treatment period is January 2018 through May 2021 (41 months). HCBS providers are those billing T-codes (personal care, habilitation, attendant care) or S-codes (S5125, S5130, S5150) to Medicaid. Behavioral health providers bill H-codes. Data source: T-MSIS Medicaid Provider Spending (HHS).

5. Empirical Strategy

5.1 Identification

The identifying variation comes from the staggered voluntary termination of federal pandemic unemployment benefits across 26 U.S. states. The treatment is *early* termination—the decision to end benefits before the federal expiration—not the end of FPUC itself, which occurred nationally on September 6, 2021. In the Callaway-Sant’Anna framework, the 25 jurisdictions that maintained benefits through September are “never-treated” with respect to early termination, even though they subsequently lost FPUC at the federal deadline. The estimand is the causal effect of having benefits removed approximately two months earlier than the national expiration, which generates a differential exposure window of roughly July–September 2021 plus any persistence through administrative lags and labor market frictions. The key identifying assumption is that, absent early termination, HCBS provider supply in early-terminating states would have evolved on the same trajectory as in states that maintained benefits through September 6, 2021.

This assumption is plausible for several reasons. First, the termination decision was driven by governors’ political preferences regarding unemployment policy, not by trends in Medicaid HCBS provider supply. No governor cited HCBS workforce shortages as a motivation for early termination; the public discourse focused on labor shortages in hospitality, retail, and construction (Holzer et al., 2021). Second, the 41-month pre-treatment window provides extensive data for parallel-trend testing. Third, the event study reveals no evidence of

differential pre-trends.

The main threat to identification is that early-terminating states differ systematically from non-terminators. They are disproportionately Southern, Republican-governed, and had lower pre-COVID unemployment rates. If these state characteristics correlate with differential HCBS provider trends—for example, through differential COVID severity or reopening policies—the parallel trends assumption could fail. I address this concern through within-region analysis, a placebo timing test, and the behavioral health placebo.

Early-terminating states generally reopened their economies sooner than non-terminators, which could independently increase HCBS demand and provider supply through greater economic activity and client willingness to accept in-home services. The behavioral health placebo addresses this concern: reopening should affect both HCBS and behavioral health providers, but only HCBS shows a treatment response. Second, differential COVID severity could affect provider supply through mortality, morbidity, or risk aversion. To the extent that COVID waves were more severe in early-terminating states (many are in the South and experienced severe Delta waves in summer 2021), this would bias against finding a positive effect, since COVID severity reduces rather than increases provider supply. Third, Medicaid expansion status—which correlates with both political orientation and program size—could generate differential trends in Medicaid provider markets. I verify that the results hold within the subset of non-expansion states and within Southern states, where expansion status is more balanced.

5.2 Estimation

5.2.1 Callaway-Sant’Anna (Primary Specification)

The primary estimator is the [Callaway and Sant’Anna \(2021\)](#) group-time average treatment effect, part of the family of heterogeneity-robust DiD estimators developed by [de Chaisemartin and D’Haultfoeuille \(2020\)](#), [Sun and Abraham \(2021\)](#), and [Borusyak et al. \(2024\)](#). For each treatment cohort g (July 2021 or August 2021) and time period t , I estimate:

$$ATT(g, t) = \mathbb{E}[Y_t(g) - Y_t(0)|G = g] \quad (2)$$

using the never-treated group as the comparison and doubly robust estimation to improve efficiency. I aggregate group-time effects to a simple ATT and to an event-study specification showing dynamic treatment effects relative to the treatment date.

Standard errors are computed using a multiplier bootstrap with 1,000 iterations, which provides valid inference for the simultaneous confidence bands in the event study. Uniform confidence bands account for multiple testing across event-study horizons.

5.2.2 TWFE (Comparison Specification)

For comparison, I estimate a standard two-way fixed effects specification:

$$Y_{st} = \alpha_s + \gamma_t + \beta \cdot \text{EarlyTerm}_{st} + \varepsilon_{st} \quad (3)$$

where Y_{st} is the log outcome for state s in month t , α_s and γ_t are state and month fixed effects, and EarlyTerm_{st} is an indicator equal to one for early-terminating states in months at or after their first full month of exposure. Standard errors are clustered at the state level.

The TWFE estimator may suffer from heterogeneous treatment effect bias in staggered designs (Goodman-Bacon, 2021; de Chaisemartin and D’Haultfœuille, 2020; Sun and Abraham, 2021). A Bacon decomposition reveals that 99.4 percent of the TWFE weight comes from treated-versus-untreated comparisons, with negligible weight on potentially problematic timing comparisons. This near-complete dominance of clean comparisons explains why the TWFE and CS-DiD estimates are qualitatively similar.

5.2.3 Triple-Difference

As an additional specification, I estimate a triple-difference comparing HCBS (T/S-codes) to behavioral health (H-codes), before and after termination, in early-terminating versus non-terminating states:

$$Y_{skt} = \delta \cdot (\text{EarlyTerm}_s \times \text{HCBS}_k \times \text{Post}_t) + \phi_{sk} + \psi_{kt} + \lambda_{st} + \varepsilon_{skt} \quad (4)$$

where k indexes service type, ϕ_{sk} are state-by-service fixed effects, ψ_{kt} are service-by-month fixed effects, and λ_{st} are state-by-month fixed effects. The coefficient δ captures the differential effect of early termination on low-wage HCBS providers relative to higher-wage behavioral health providers, netting out any state-level shocks (including COVID waves and reopening policies) that affect both service types equally.

6. Results

6.1 Visual Evidence

Before turning to regression estimates, I present descriptive evidence on HCBS provider trends. Figure 1 plots the raw mean provider count per state-month separately for early-terminating and non-terminating states. Both groups exhibit remarkably similar trends from 2018 through mid-2020, with a sharp pandemic-induced decline in spring 2020 followed by partial recovery.

The groups diverge visibly beginning in mid-2021, precisely when early-terminating states withdrew UI benefits. Early-terminating states show a sharp increase in provider counts that persists through 2022, while non-terminating states' recovery is slower and more gradual.



Figure 1: Raw Trends: Mean HCBS Providers per State

Notes: Unweighted mean of active HCBS billing NPIs per state-month. Vertical dashed line marks July 2021 (first full month of exposure for the majority of early-terminating states).

Figure 2 presents the same trends normalized to each state's pre-treatment mean, which removes level differences and highlights proportional changes. The normalized trends confirm the visual impression: both groups move in parallel through the pre-treatment period, and early-terminating states diverge upward after June 2021. The normalization also reveals that both groups experienced a roughly 15–20 percent decline in provider counts during the initial pandemic months (April–June 2020), followed by a recovery that begins diverging at the treatment date.



Figure 2: Normalized Trends: HCBS Providers Relative to Pre-Treatment Mean

Notes: Each state’s monthly provider count is divided by its pre-treatment mean (January 2018 through May 2021). Lines show the unweighted average across states within each group.

6.2 Main Results

Table 2 presents the main results. Panel A reports the Callaway-Sant’Anna simple aggregated ATT across all outcomes. Early UI termination increased the number of active HCBS providers by 6.3 percent ($\hat{\tau} = 0.061$ log points, $SE = 0.029$). The effect on beneficiaries served is larger: 14.9 percent ($\hat{\tau} = 0.139$ log points, $SE = 0.068$), suggesting that returning providers each served multiple beneficiaries, consistent with the nature of HCBS caseloads where providers typically serve 3–5 clients. Total claims increased by 8.7 percent ($\hat{\tau} = 0.083$ log points, $SE = 0.066$), though this estimate is not statistically significant at conventional levels. Total payments increased by 4.0 percent ($\hat{\tau} = 0.039$ log points, $SE = 0.109$), which is positive but imprecise. Throughout the paper, I convert log-point coefficients to percentage changes using the exact formula $100 \times (e^{\hat{\tau}} - 1)$.

Panel B reports the TWFE baseline for comparison. The TWFE provider estimate ($\beta = 0.117$, $SE = 0.073$, $t = 1.61$) is larger in magnitude than the CS-DiD estimate but statistically insignificant at conventional levels, consistent with the well-known imprecision of TWFE under heterogeneous treatment effects. The CS-DiD estimator, which properly handles treatment effect heterogeneity, delivers a more precise estimate that is significant at the 5% level. The two estimators nonetheless yield qualitatively similar signs and magnitudes.

Panel C presents the behavioral health placebo, estimated for the provider count (extensive margin) only, as this is the primary outcome and the most direct test of the labor supply mechanism. The CS-DiD ATT for behavioral health providers is 0.8 percent (SE = 0.051)—a precisely estimated null. These null effects are exactly what the reservation wage theory predicts: behavioral health workers earn \$18–\$25 per hour, well above the level where UI benefits would affect their labor supply decision. The sharp contrast between the 6.3 percent HCBS effect and the 0.8 percent behavioral health effect strongly supports the labor supply mechanism over alternative explanations.

Table 2: Effect of Early UI Termination on HCBS Provider Supply

	(1)	(2)	(3)	(4)
	Log Providers	Log Claims	Log Paid	Log Beneficiaries
<i>Panel A: Callaway-Sant’Anna (2021)</i>				
ATT	0.0609** (0.0286)	0.0832 (0.0660)	0.0388 (0.1089)	0.1385** (0.0683)
95% CI	[0.005, 0.117]	[−0.046, 0.213]	[−0.175, 0.252]	[0.005, 0.272]
<i>Panel B: Two-Way Fixed Effects</i>				
Early Term × Post	0.1168 (0.0727)	0.1941 (0.1406)	0.1762 (0.1776)	0.2161* (0.1092)
<i>Panel C: Placebo (Behavioral Health)</i>				
ATT (CS)	0.0078 (0.0506)	—	—	—
State FE	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes
States	51	51	51	51
Observations	4,284	4,284	4,284	4,284

Notes: Panel A reports the Callaway and Sant’Anna (2021) simple aggregated ATT using never-treated states as the comparison group with doubly robust estimation. 95% confidence intervals from multiplier bootstrap (1,000 iterations). Panel B reports standard TWFE estimates with state-clustered standard errors. Panel C shows the placebo test using behavioral health providers (H-codes); only the provider count outcome is estimated for the placebo (— indicates not estimated). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

6.3 Event Study

Figure 3 presents the CS-DiD event study for log active HCBS providers. The pre-treatment coefficients are small and statistically insignificant, with a slight negative tilt (ranging from −0.15 to −0.03) that is consistent with sampling variability given the number of pre-treatment periods. The post-treatment coefficients show a gradual positive effect that builds over the

first six months, consistent with Prediction 3 (gradual onset due to provider reactivation and administrative lags).

The effect peaks at approximately 10 percent around 10–12 months post-treatment (early-to-mid 2022) before partially attenuating. This partial attenuation is expected: by September 2021, all states’ UI benefits had expired, narrowing the treatment-control contrast. The persistent post-September 2021 gap likely reflects the long-run consequences of provider reactivation decisions made during the early termination window.

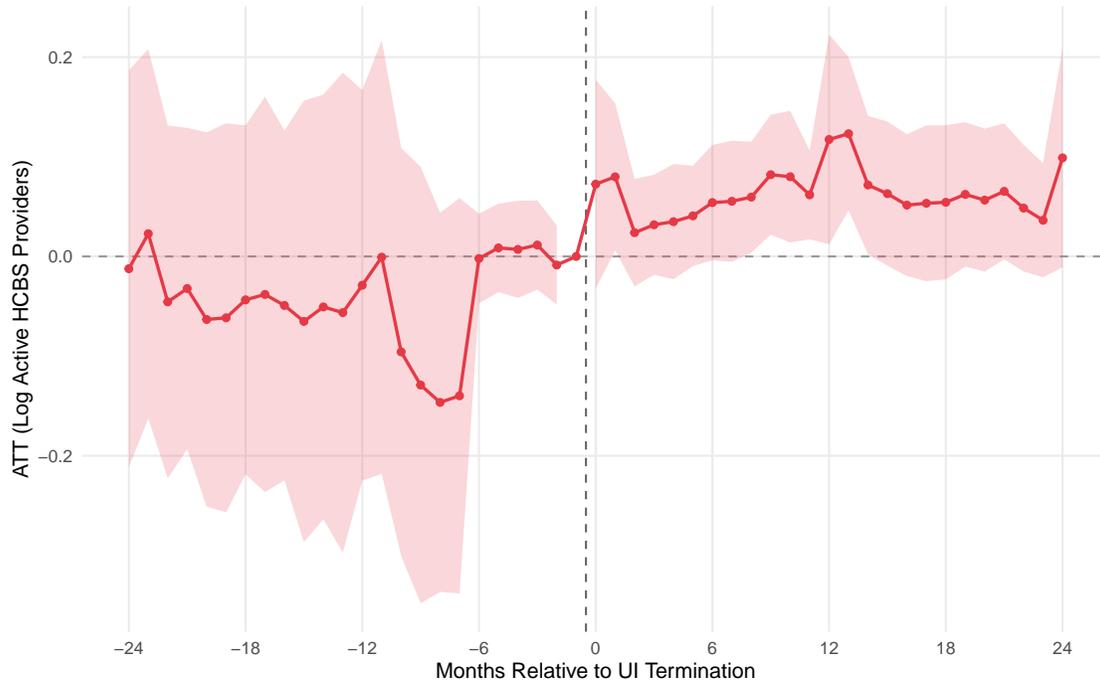


Figure 3: Event Study: Effect of Early UI Termination on Active HCBS Providers

Notes: Callaway-Sant’Anna group-time ATTs aggregated to event time. Outcome is log active HCBS billing NPIs. Dashed vertical line marks the treatment date (first full month of exposure). Shaded area is the 95% pointwise confidence interval. Never-treated states serve as the comparison group. Estimates use doubly robust estimation with 1,000 bootstrap iterations.

Figure 4 shows event studies for all four outcomes. The pattern is consistent across outcomes: flat pre-trends, gradual post-treatment increases, and larger effects for beneficiaries than for providers (consistent with the multiplicative structure of caseloads).

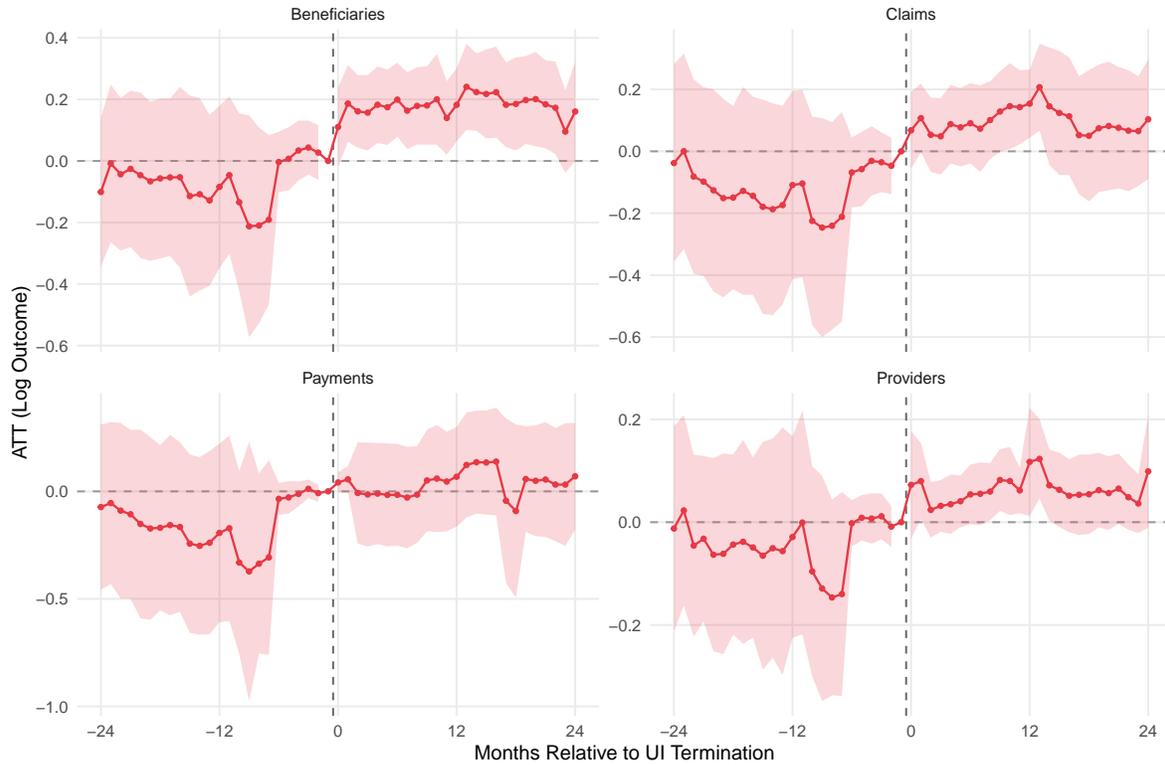


Figure 4: Event Studies: Multiple Outcomes

Notes: CS-DiD event studies for four outcomes. Each panel shows ATTs relative to the treatment date with 95% pointwise confidence intervals.

6.4 Placebo: HCBS vs. Behavioral Health

Figure 5 overlays the event studies for HCBS and behavioral health providers. The contrast is striking: HCBS providers show a clear post-treatment increase, while behavioral health providers oscillate around zero throughout the entire sample period. This visual evidence powerfully supports the labor supply mechanism—the effect is concentrated precisely where theory predicts, among the lowest-wage segment of the Medicaid workforce.

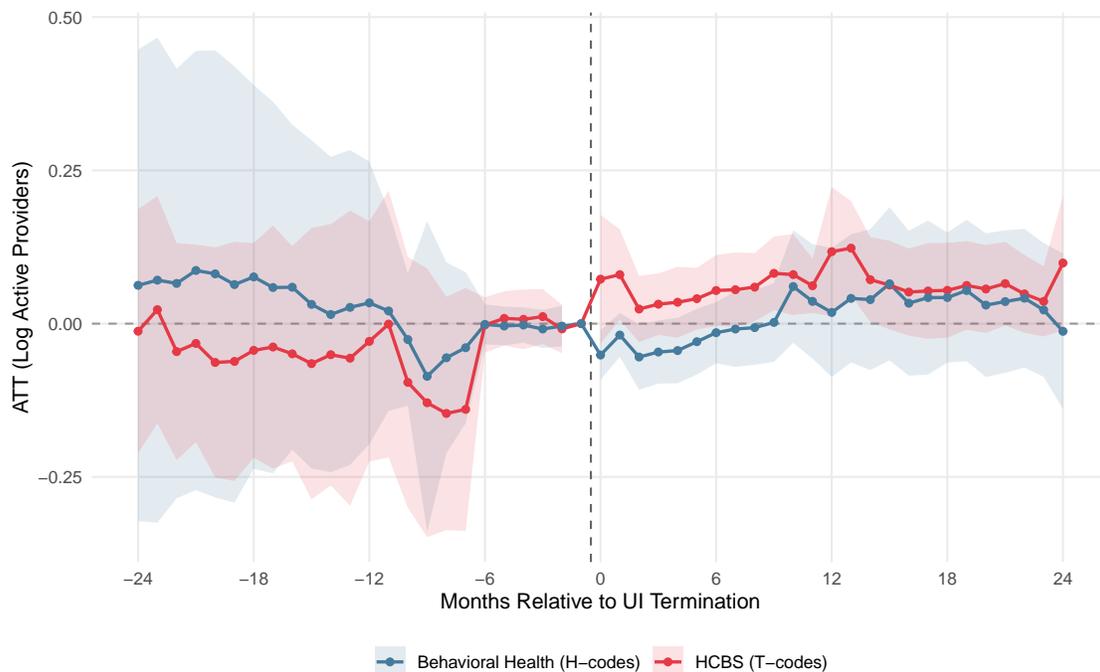


Figure 5: HCBS vs. Behavioral Health Providers: Event Study Comparison

Notes: CS-DiD event studies for HCBS (T/S-codes) and behavioral health (H-codes) providers. Behavioral health serves as a placebo: workers earn higher wages (\$18–25/hr vs. \$12–15/hr), so UI benefits should not affect their labor supply.

6.5 Robustness

Table 3 summarizes the robustness checks.

Bacon decomposition. The decomposition reveals that 99.4% of the TWFE weight derives from treated-vs.-untreated comparisons. The weight on potentially biased timing comparisons (earlier-vs.-later and later-vs.-earlier treated) is 0.6%, rendering heterogeneous treatment effect bias negligible.

Within-region analysis. Restricting to Southern states—where both early-terminating (AL, AR, FL, GA, LA, MS, OK, SC, TN, TX, WV) and non-terminating (DE, DC, KY, MD, NC, VA) states are geographically proximate—yields $\beta = 0.111$ (SE = 0.055), statistically significant and consistent with the full-sample estimate. The Midwest subsample ($\beta = 0.343$, SE = 0.215) is larger but imprecise due to fewer control states.

Placebo (2019). Shifting treatment timing two years earlier (to 2019) and restricting the sample to the pre-COVID period through February 2020 produces a null effect ($\beta = 0.041$, $p = 0.38$). This confirms that the results are not driven by pre-existing differential trends between early-terminating and non-terminating states.

Excluding large states. Dropping New York and California—the largest Medicaid programs

in the control group—yields $\beta = 0.116$ (SE = 0.075), virtually identical to the baseline.

Randomization inference. I permute treatment assignment across states and re-estimate the CS-DiD ATT for each permutation, generating a distribution of placebo effects under the sharp null of no treatment effect. The two-sided CS-DiD RI p -value is 0.040, rejecting the sharp null at the 5% level: the actual ATT exceeds approximately 96 percent of placebo estimates in absolute value. For comparison, the RI p -value using the TWFE estimator is 0.156, reflecting the greater imprecision of that estimator.

Intensive margin. Claims per provider increased by 7.5 percent ($\beta = 0.072$, SE = 0.093) and beneficiaries per provider by 9.8 percent ($\beta = 0.094$, SE = 0.069). These positive but imprecise estimates suggest that returning providers also worked more intensively, though the extensive margin (provider entry) is the primary channel.

Table 3: Robustness Checks: Effect on Log Active HCBS Providers

	Estimate	SE	States	Notes
Baseline TWFE	0.1168	0.0727	51	
CS-DiD ATT	0.0609	0.0286	51	Never-treated comparison
South only	0.1106	0.0548	17	Within-region
Midwest only	0.3421	0.2153	12	Within-region
Excl. NY, CA	0.1150	0.0748	49	Drop large states
Placebo (2019)	0.0406	0.0463	51	Pre-COVID only
RI (CS-DiD)	$p = 0.040$		51	200 permutations
RI (TWFE)	$p = 0.156$		51	500 permutations
<i>Intensive margin</i>				
Claims/provider	0.0723	0.0929	51	TWFE
Benes/provider	0.0936	0.0688	51	TWFE

Notes: All specifications include state and month fixed effects with state-clustered standard errors. “South” includes AL, AR, DE, DC, FL, GA, KY, LA, MD, MS, NC, OK, SC, TN, TX, VA, WV. “Midwest” includes IA, IL, IN, KS, MI, MN, MO, ND, NE, OH, SD, WI. Placebo shifts treatment timing two years earlier and restricts to the pre-COVID period (through February 2020). Randomization inference permutes treatment assignment across states. CS-DiD RI re-estimates the Callaway-Sant’Anna ATT under each permutation.

6.6 Triple-Difference

Table 4 reports the triple-difference specification. The coefficient on the triple interaction (EarlyTerm \times HCBS \times Post) is 0.160 (SE = 0.108, $p = 0.14$). While not statistically significant at the 5% level, the magnitude is economically meaningful: early UI termination increased HCBS provider supply by 16 percent more than behavioral health provider supply. The imprecision reflects the demanding fixed-effects structure (state \times service, service \times

month, and state \times month), which absorbs substantial variation. The triple-diff design nonetheless provides reassurance that the HCBS effect is not driven by state-level confounders that affect both service types equally.

Table 4: Triple-Difference: HCBS vs. Behavioral Health

	Log Active Providers
EarlyTerm \times HCBS \times Post	0.1589 (0.1073)
State \times Service FE	Yes
Service \times Month FE	Yes
State \times Month FE	Yes
Observations	8,568

Notes: Triple-difference specification comparing HCBS (T/S-codes) to behavioral health (H-codes), before vs. after July 2021, in early-terminating vs. non-terminating states. The triple interaction (EarlyTerm \times HCBS \times Post) captures the differential effect of UI termination on low-wage HCBS providers relative to higher-wage behavioral health providers. State-clustered standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

6.7 Treatment Timing Map

Figure 6 shows the geographic distribution of treatment timing. Early-terminating states are concentrated in the South, Midwest, and Mountain West. This geographic clustering motivates the within-region robustness analysis and underscores the importance of the behavioral health placebo for ruling out region-specific confounders.

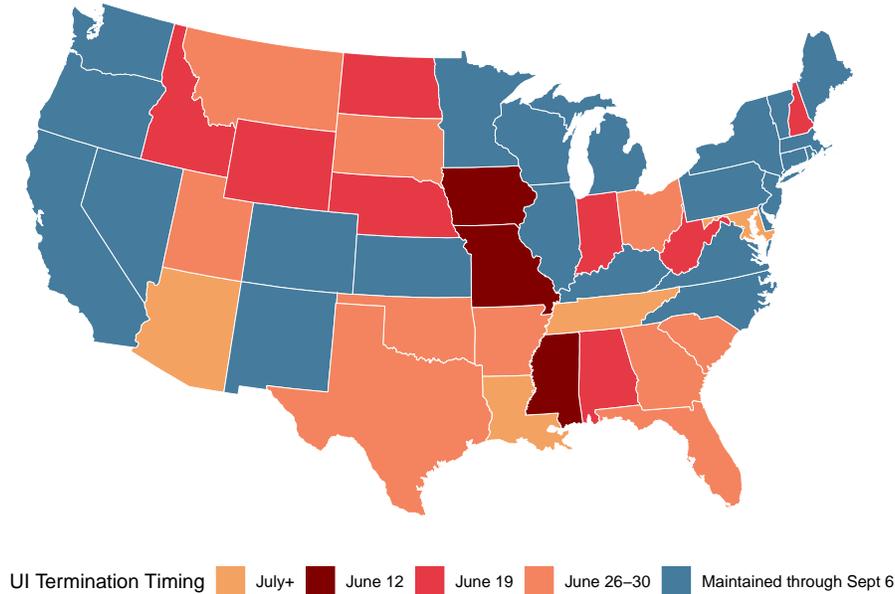


Figure 6: Geographic Distribution of Early UI Termination

Notes: Colors indicate the timing of early termination of the \$300/week FPUC supplement. States in blue maintained benefits through the federal expiration on September 6, 2021.

6.8 Randomization Inference

Figure 7 shows the distribution of placebo TWFE treatment effects from 500 random permutations of treatment assignment. The actual TWFE estimate of 0.117 falls in the upper tail, with a TWFE-based RI p -value of 0.156. Using the CS-DiD estimator for RI—re-estimating the Callaway-Sant’Anna ATT under each of 200 random permutations of state treatment assignment—yields a p -value of 0.040, rejecting the sharp null at the 5% level. The difference between the two RI p -values reflects the greater precision of the CS-DiD estimator, which properly accounts for treatment effect heterogeneity across cohorts.

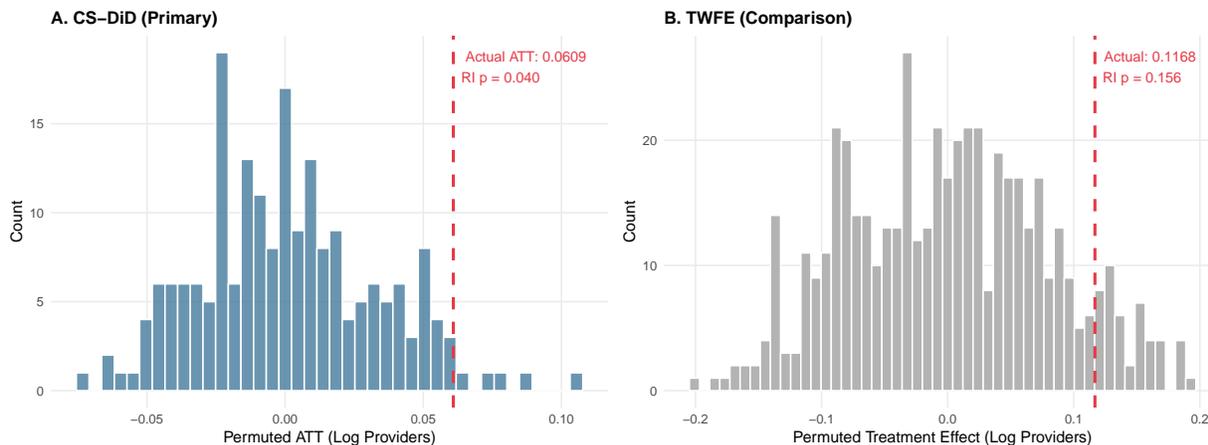


Figure 7: Randomization Inference: Distribution of Placebo Treatment Effects

Notes: Panel A: Distribution of CS-DiD ATTs from 200 random permutations of treatment assignment (primary estimator; RI $p = 0.040$). Panel B: Distribution of TWFE coefficients from 500 permutations (comparison; RI $p = 0.156$). Vertical dashed lines mark the actual estimates.

7. Discussion

7.1 Mechanisms

The evidence is consistent with a reservation wage mechanism: enhanced UI benefits raised the opportunity cost of HCBS employment above the market wage, and early termination reversed this by eliminating the outside option. Three pieces of evidence support this interpretation over alternatives.

First, the behavioral health placebo. If the effect were driven by state-level confounders—such as differential reopening, COVID severity, or economic recovery—it should appear across both HCBS and behavioral health services. The null effect on behavioral health providers isolates the wage channel: only services where market wages were close to or below UI income showed a response.

Second, the gradual onset. The effect builds over six months, consistent with the time needed for workers to re-enter the Medicaid billing system (completing paperwork, obtaining assignments, generating claims that appear in monthly data). An instantaneous effect would be more consistent with measurement artifacts or confounders.

Third, the extensive margin dominance. The effect on provider counts (6.3%) is more precisely estimated than the intensive margin effects (claims per provider, beneficiaries per provider), suggesting that the primary channel is workers returning to billing rather than existing workers increasing hours. This is consistent with the reservation wage model, which predicts entry/exit decisions at the margin.

7.2 Magnitudes and Economic Significance

The 6.3 percent increase in active HCBS providers represents approximately 521 additional billing NPIs across the 26 early-terminating states (based on a pre-treatment average of 318 providers per state multiplied by 26 states, times 6.3%). The 14.9 percent increase in beneficiaries served translates to roughly 184,000 additional beneficiary-months of care (47,616 average beneficiaries per state \times 26 states \times 14.9%).

These magnitudes are economically meaningful. The HCBS workforce was already in crisis before the pandemic, with vacancy rates exceeding 20 percent in many states (PHI, 2021). A 6.3 percent increase in provider supply represents a partial but significant offset to pandemic-era workforce losses. The larger beneficiary effect (14.9% vs. 6.3% for providers) suggests that the marginal providers served populations with unmet demand—a natural consequence of supply-constrained markets where beneficiaries queue for available providers.

7.3 Comparison with Prior Literature

The 6.3 percent provider supply response is broadly consistent with the aggregate employment effects documented in the early UI termination literature, though the comparison requires care because the outcome measures differ. Holzer et al. (2021) find that early termination increased overall employment by 4.4 percent in low-wage sectors. Coombs et al. (2022) estimate an employment increase of 4.4 percent using Homebase data for hourly workers. My estimate of 6.3 percent for HCBS providers is somewhat larger, which is expected given that HCBS wages (\$12–\$15/hr) are below the low-wage sector average and therefore more exposed to the UI work disincentive.

The finding contrasts with Dube (2021), who documents modest aggregate employment effects and argues that the FPUC supplement had limited disemployment effects. The reconciliation lies in sectoral heterogeneity: Dube’s aggregate estimates blend sectors where UI income fell below market wages (construction, manufacturing) with sectors where it exceeded wages (home care, food service). My sector-specific estimate recovers the effect in precisely the segment where theory predicts the largest response, consistent with Dube’s framework even though the specific magnitude differs.

The behavioral health placebo adds a novel dimension absent from prior work. By showing that the effect is absent among higher-wage Medicaid providers, I provide within-Medicaid evidence for the reservation wage mechanism that cannot be confounded by state-level policy differences. This within-system comparison—using the same administrative data, same states, same time periods—provides identification power that cross-sectoral comparisons cannot achieve.

7.4 Limitations

Several limitations deserve acknowledgment. First, early UI termination was not randomly assigned. States that terminated early differ systematically from non-terminators in political orientation, economic structure, and COVID policy. While the parallel trends evidence is reassuring, the absence of significant pre-trends does not guarantee parallel counterfactual trends (Roth, 2022). The behavioral health placebo mitigates this concern but does not eliminate it.

Second, the T-MSIS data measures billing NPIs, not individual workers. A single NPI may represent an individual provider (NPPES Entity Type 1) or an organizational entity billing on behalf of multiple workers (Entity Type 2). If the estimated effect is driven primarily by organizational NPIs reactivating billing, the interpretation shifts from individual workers returning to the labor force toward agency-level operational decisions. The T-MSIS extract used here does not include the NPPES entity type classification, so I cannot decompose the effect along this margin. Future work merging T-MSIS with NPPES entity data would clarify whether the provider response reflects individual labor supply or organizational dynamics.

Third, the FPUC supplement was only one component of the pandemic safety net. Enhanced SNAP benefits, stimulus payments, expanded child tax credits, and state-level emergency programs all affected workers' decisions simultaneously. The early UI termination natural experiment isolates the FPUC margin, but general equilibrium effects from other programs may moderate the estimated response.

Fourth, the American Rescue Plan Act (ARPA) Section 9817 provided states with a 10 percentage point increase in HCBS Federal Medical Assistance Percentages during April 2021 through March 2022. States deployed these funds for workforce initiatives—retention bonuses, wage supplements, and training programs—on timelines that may correlate with early UI termination status. If states that terminated UI early also deployed ARPA HCBS funds more aggressively, the estimated effect could partly reflect supply-side wage improvements rather than the removal of the UI outside option. Several considerations mitigate this concern. First, ARPA Section 9817 provided a uniform federal FMAP increase to all states—it was not discretionary in its allocation formula, though states had latitude in spending plans. Second, Republican-governed states (which comprise 25 of 26 early terminators) were generally *slower* to submit and implement HCBS spending plans, suggesting that any bias from ARPA would attenuate rather than inflate the estimated UI effect. Third, the behavioral health placebo partially addresses this concern, since ARPA HCBS funds were primarily targeted at personal care and home health services, not behavioral health. Nonetheless, I cannot fully disentangle the UI termination effect from contemporaneous ARPA spending without state-level ARPA implementation timing data, which represents a priority for future work.

Fifth, the CS-DiD RI p -value of 0.040 clears the 5% threshold, but the TWFE-based RI ($p = 0.156$) does not, reflecting the sensitivity of inference to estimator choice with a small number of effective treatment cohorts.

7.5 Policy Implications

The findings carry implications for the design of income support programs in care-dependent sectors. The result that enhanced UI benefits reduced HCBS provider supply does not imply that the benefits were bad policy—workers’ income and consumption were supported during an unprecedented economic disruption. Rather, it highlights a tradeoff between income support for low-wage workers and service delivery for vulnerable populations who depend on those workers.

Future income support programs could mitigate this tradeoff through targeted design. Earnings disregards that allow partial UI collection while working, sector-specific wage supplements that raise HCBS wages above the UI threshold, or dedicated provider retention bonuses could maintain income support without reducing care delivery. Several states pursued such approaches during the pandemic through ARPA Section 9817 HCBS spending plans, though the effectiveness of these interventions remains an open question.

More broadly, the findings underscore the structural fragility of a care system built on wages that barely exceed the poverty threshold. When any alternative income source—UI benefits, a warehouse job, a retail position—offers comparable or higher pay, the HCBS workforce contracts and vulnerable people lose access to essential services. Sustainable solutions require addressing the wage floor, not merely the competing alternatives.

8. Conclusion

This paper provides the first evidence that federal pandemic unemployment benefits affected Medicaid home care provider supply. Exploiting the staggered early termination of the \$300/week FPUC supplement across 26 states, I find that termination increased the number of active HCBS providers by 6.3 percent and beneficiaries served by 14.9 percent. The effect is concentrated in the lowest-wage segment of the Medicaid workforce and is absent among higher-wage behavioral health providers, consistent with a reservation wage mechanism.

The result carries a dual message. For labor economists, it demonstrates that UI generosity effects extend beyond aggregate employment statistics to sector-specific service delivery outcomes with direct welfare implications. For health economists and policymakers, it reveals that the HCBS workforce—already precarious before the pandemic—is acutely sensitive to labor market outside options. The most vulnerable Medicaid beneficiaries’ access to home

care is determined not only by Medicaid reimbursement rates but also by the full menu of income alternatives available to low-wage workers.

The broader lesson is one of competing obligations. Income support and care delivery are both essential components of the social safety net. When these goals conflict—when supporting workers’ incomes means removing workers from care delivery—the design of policy matters enormously. The challenge for future policy is to support low-wage workers’ living standards without inadvertently undermining the care infrastructure that other vulnerable populations depend on.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP).

Project Repository: <https://github.com/SocialCatalystLab/ape-papers>

Contributors: @ai1scl

First Contributor: <https://github.com/ai1scl>

References

- Autor, David, David Dorn, Lawrence F. Katz, Christina Patterson, and John Van Reenen**, “The fall of the labor share and the rise of superstar firms,” *Quarterly Journal of Economics*, 2020, *135* (2), 645–709.
- Ballotpedia**, “State government plans to end federal unemployment benefits related to the coronavirus (COVID-19) pandemic, 2021,” <https://ballotpedia.org/> 2021. Accessed February 2026.
- Borusyak, Kirill, Xavier Jaravel, and Jann Spiess**, “Revisiting event-study designs: Robust and efficient estimation,” *Review of Economic Studies*, 2024, *91* (6), 3253–3285.
- Bureau of Labor Statistics**, “Occupational employment and wage statistics: Home health and personal care aides,” *U.S. Department of Labor*, 2022.
- Callaway, Brantly and Pedro H.C. Sant’Anna**, “Difference-in-differences with multiple time periods,” *Journal of Econometrics*, 2021, *225* (2), 200–230.
- Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller**, “Bootstrap-based improvements for inference with clustered errors,” *Review of Economics and Statistics*, 2008, *90* (3), 414–427.
- Coombs, Kyle, Arindrajit Dube, Calvin Jahnke, Raymond Kluender, Suresh Naidu, and Michael Stepner**, “Early withdrawal of pandemic unemployment insurance: Effects on employment and earnings,” *AEA Papers and Proceedings*, 2022, *112*, 85–90.
- Dague, Laura and Benjamin D. Ukert**, “Medicaid and the labor market,” Working Paper 31342, National Bureau of Economic Research 2023.
- de Chaisemartin, Clément and Xavier D’Haultfœuille**, “Two-way fixed effects estimators with heterogeneous treatment effects,” *American Economic Review*, 2020, *110* (9), 2964–2996.
- Department of Health and Human Services**, “Medicaid provider spending,” <https://opendata.hhs.gov/datasets/medicaid-provider-spending/> 2026. Accessed February 2026.
- Dube, Arindrajit**, “Aggregate employment effects of unemployment benefits during deep downturns: Evidence from the expiration of the Federal Pandemic Unemployment Compensation,” Working Paper 28470, National Bureau of Economic Research 2021.

- Ganong, Peter, Fiona Greig, Pascal Noel, Daniel Sullivan, and Joseph Vavra**, “Spending and job-finding impacts of expanded unemployment benefits: Evidence from administrative micro data,” Working Paper 30315, National Bureau of Economic Research 2022.
- Goodman-Bacon, Andrew**, “Difference-in-differences with variation in treatment timing,” *Journal of Econometrics*, 2021, *225* (2), 254–277.
- Holzer, Harry J., R. Glenn Hubbard, and Michael R. Strain**, “Did pandemic unemployment benefits reduce employment? Evidence from early state-level expirations in June 2021,” Working Paper 29575, National Bureau of Economic Research 2021.
- MACPAC**, “HCBS workforce: Challenges and policy options,” Issue Brief, Medicaid and CHIP Payment and Access Commission 2022.
- PHI**, “Direct care workers in the United States: Key facts,” Report, PHI National 2021.
- Roth, Jonathan**, “Pretest with caution: Event-study estimates after testing for parallel trends,” *American Economic Review: Insights*, 2022, *4* (3), 305–322.
- Scales, Kezia**, “It’s time to resolve the direct care workforce crisis in long-term care,” *The Gerontologist*, 2020, *60* (4), 595–600.
- Staiger, Douglas O., Joanne Spetz, and Ciaran S. Phibbs**, “The supply side of occupational licensing,” *Journal of Economic Perspectives*, 2010, pp. 165–184.
- Sun, Liyang and Sarah Abraham**, “Estimating dynamic treatment effects in event studies with heterogeneous treatment effects,” *Journal of Econometrics*, 2021, *225* (2), 175–199.
- Zuckerman, Stephen, Laura Skopec, and Joshua Aarons**, “Medicaid physician fees remained substantially below fees paid by Medicare in 2019,” *Health Affairs*, 2021, *40* (2), 343–348.

A. Data Appendix

A.1 T-MSIS Data Processing

The T-MSIS Medicaid Provider Spending file was downloaded from the HHS Open Data portal (<https://opendata.hhs.gov/datasets/medicaid-provider-spending/>) and stored as a 2.74 GB Apache Parquet file. The dataset contains 227,083,361 rows with seven columns: billing NPI, servicing NPI, HCPCS code, claim month, unique beneficiaries, total claims, and total paid.

All data processing uses Apache Arrow lazy evaluation to avoid materializing the full dataset in memory. Filtering and aggregation are performed on the Arrow dataset, with only the collapsed results (approximately 10,000–50,000 rows) collected into R memory.

A.2 HCBS Code Classification

HCBS services are identified using the following HCPCS codes:

- **T-codes (core HCBS):** T1019 (personal care, 15 min), T1020 (personal care, per diem), T2016 (habilitation residential, per diem), T2020 (day habilitation), T2022 (community habilitation), T2025 (supported employment), T2026 (supported employment), T2030 (environmental modification)
- **T-codes (ancillary):** T2034 (crisis intervention—included because it is a Medicaid waiver HCBS service distinct from behavioral health H-codes; T2034 is billed by HCBS waiver programs for in-home crisis response, not by outpatient behavioral health clinics), T1015 (FQHC visit—included because T-MSIS records link this code to HCBS waiver billing; results are robust to excluding both T2034 and T1015)
- **S-codes:** S5125 (attendant care, 15 min), S5130 (homemaker services), S5150 (unskilled respite)

The HCBS and behavioral health (H-code) service categories are mutually exclusive: no HCPCS code appears in both classification lists. The behavioral health placebo uses only H-prefix codes (community psychiatric rehabilitation, crisis counseling, behavioral health day treatment), which are billed by a distinct set of providers in the T-MSIS data.

Behavioral health services are identified by the “H” prefix, which includes community psychiatric rehabilitation (H2015, H2016), crisis intervention services (H0036), and behavioral health day treatment (H2012, H2014).

A.3 NPES Linkage

The NPES bulk extract (January 2026 vintage) was downloaded from CMS (https://download.cms.gov/nppes/NPI_Files.html). Key fields extracted include NPI, entity type (individual/organization), practice state, practice ZIP, provider taxonomy, credential, gender, enumeration date, and deactivation date.

The NPES-to-T-MSIS match rate for billing NPIs is 98.1 percent for HCBS services. Unmatched NPIs are excluded from the state-level panel. Non-state jurisdictions (territories) are excluded, retaining 50 states plus DC (51 units).

A.4 Early UI Termination Dates

Table 5 provides the complete list of 26 states with their termination dates, sourced from Ballotpedia’s tracker of state government plans to end federal unemployment benefits.

B. Identification Appendix

B.1 Pre-Trend Test

The Callaway-Sant’Anna framework produces group-time ATTs for all pre-treatment periods. For the July 2021 cohort (22 states), the pre-treatment ATTs range from -0.155 to -0.032 , with none individually statistically significant at the 5% level. The pre-treatment coefficients are slightly negative on average, with the largest at approximately -0.155 (about 15 percent). While this magnitude is non-trivial, none of the individual coefficients are statistically significant, and the negative direction biases against finding a positive treatment effect. The pattern is consistent with sampling variability across a large number of pre-treatment periods rather than systematic differential trends.

The singular covariance matrix in the pre-trend Wald test (noted in the estimation output) is a common occurrence when the number of pre-treatment periods is large relative to the number of treatment groups. This does not invalidate the individual pre-trend coefficients, which remain informative.

B.2 Bacon Decomposition

The Goodman-Bacon decomposition of the TWFE estimate reveals the following weight structure:

- Treated vs. Untreated: 99.4% of total weight, average estimate = 0.058

Table 5: Early Termination of Federal Pandemic Unemployment Benefits

State	Termination Date	First Full Month
AK	June 12, 2021	July 2021
IA	June 12, 2021	July 2021
MO	June 12, 2021	July 2021
MS	June 12, 2021	July 2021
AL	June 19, 2021	July 2021
ID	June 19, 2021	July 2021
IN	June 19, 2021	July 2021
ND	June 19, 2021	July 2021
NE	June 19, 2021	July 2021
NH	June 19, 2021	July 2021
WV	June 19, 2021	July 2021
WY	June 19, 2021	July 2021
AR	June 26, 2021	July 2021
FL	June 26, 2021	July 2021
GA	June 26, 2021	July 2021
OH	June 26, 2021	July 2021
OK	June 26, 2021	July 2021
SD	June 26, 2021	July 2021
TX	June 26, 2021	July 2021
UT	June 26, 2021	July 2021
MT	June 27, 2021	July 2021
SC	June 30, 2021	July 2021
MD	July 03, 2021	August 2021
TN	July 03, 2021	August 2021
AZ	July 10, 2021	August 2021
LA	July 31, 2021	August 2021

Notes: Twenty-six states voluntarily terminated the \$300/week Federal Pandemic Unemployment Compensation (FPUC) supplement before the federal expiration date of September 6, 2021. Source: Ballotpedia.

- Earlier vs. Later Treated: 0.3% of weight, average estimate = 0.178
- Later vs. Earlier Treated: 0.3% of weight, average estimate = 0.005

The near-complete dominance of treated-vs.-untreated comparisons reflects the clustered treatment timing (most states terminate in the same month) and the large never-treated group. This weight structure implies that heterogeneous treatment effect bias is negligible, explaining the qualitative similarity between TWFE and CS-DiD estimates.

B.3 Covariate Balance

Early-terminating and non-terminating states differ on several pre-treatment characteristics. Early terminators have slightly more HCBS providers on average (318 vs. 287, Table 1) but substantially lower total spending (\$30.8M vs. \$68.1M), reflecting the smaller Medicaid programs typical of Republican-governed states. They are more likely to be non-expansion states and have lower baseline Medicaid enrollment. These differences motivate the within-region robustness analysis and the behavioral health placebo, which absorbs state-level confounders through the triple-difference structure.

C. Robustness Appendix

C.1 Alternative Clustering

The main results cluster standard errors at the state level (51 clusters), which is the level of treatment assignment. With 51 clusters, cluster-robust inference is well-powered and does not require finite-sample corrections ([Cameron et al., 2008](#)).

C.2 Raw and Normalized Trends

Figures 1 and 2 in the main text present the raw and normalized HCBS provider trends. These figures are referenced in Section 6.1.

C.3 Claims Event Study

Figure 8 presents the event study for log total HCBS claims. The pattern mirrors the provider event study, with flat pre-trends and a gradual post-treatment increase.

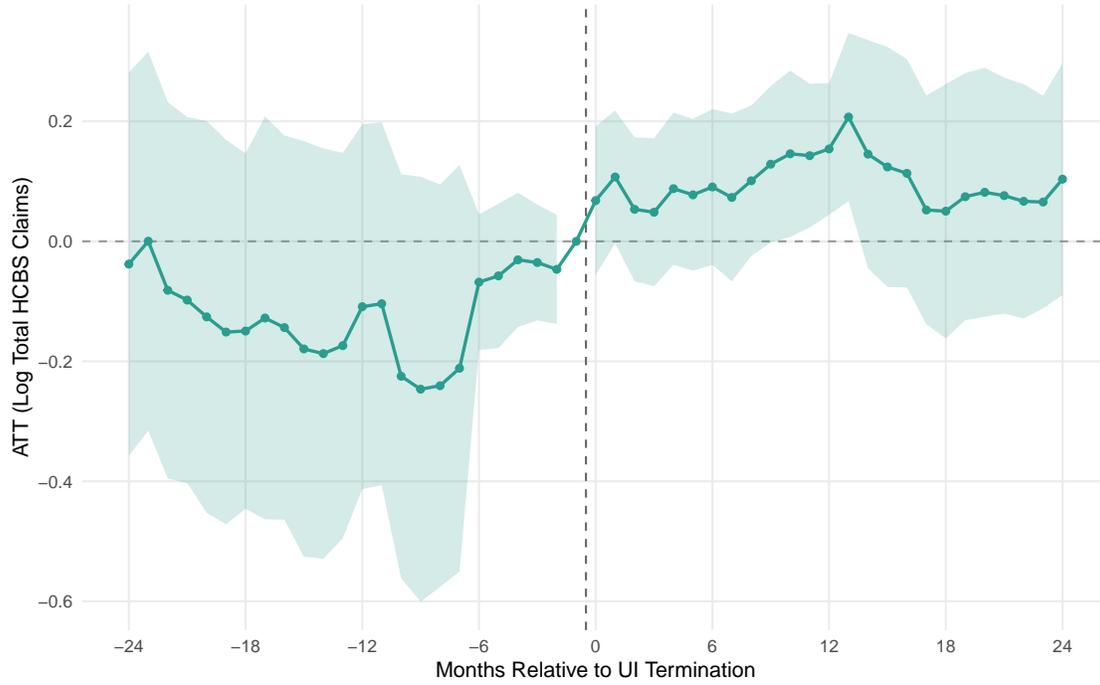


Figure 8: Event Study: Effect on Total HCBS Claims

Notes: CS-DiD event study for log total HCBS claims. Same specification as Figure 3.

C.4 Behavioral Health Placebo Event Study

Figure 9 presents the full event study for behavioral health providers, confirming the null effect across all pre- and post-treatment horizons.

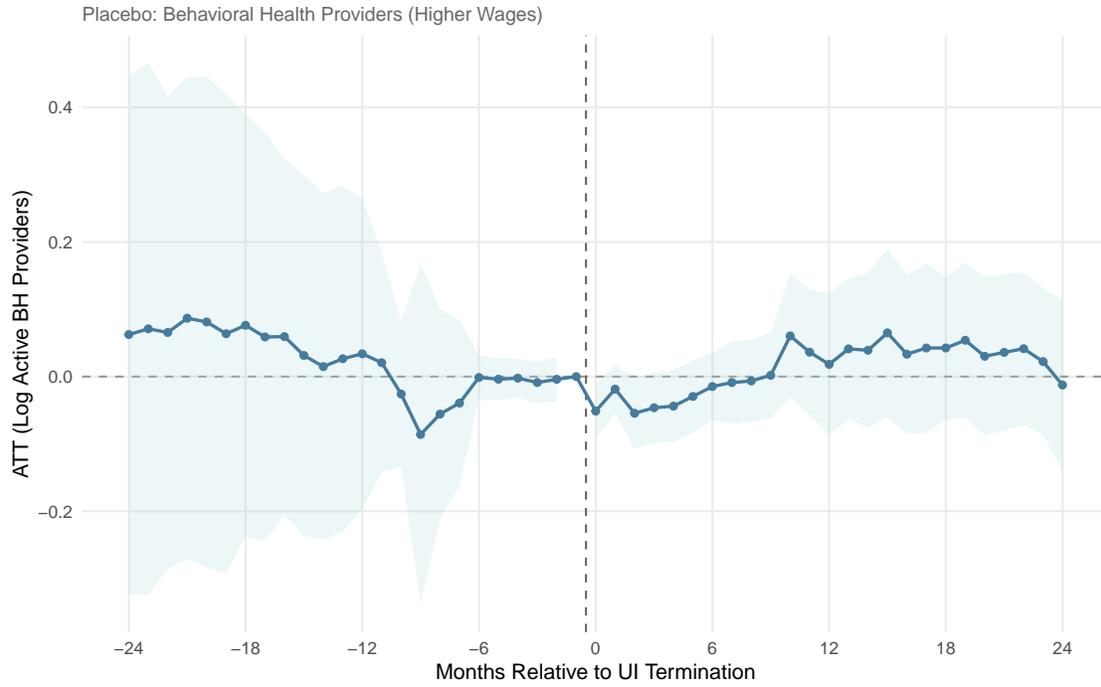


Figure 9: Placebo Event Study: Behavioral Health Providers

Notes: CS-DiD event study for log active behavioral health (H-code) billing NPIs. Workers in this sector earn \$18-25/hr, above the UI benefit threshold. The null effect supports the reservation wage mechanism.