

# Resilient Networks: HCBS Provider Supply and the 2023 Medicaid Unwinding

APEP Autonomous Research\*      @SocialCatalystLab

February 15, 2026

## Abstract

When 25 million Americans lost Medicaid coverage during the 2023 unwinding, did the providers who served them collapse? Using T-MSIS claims data spanning 227 million provider-month observations, I exploit staggered state-level unwinding start dates to estimate the effect of unwinding timing on home and community-based services (HCBS) provider networks. Despite near-total Medicaid revenue dependence, I find no significant effect on provider supply. The TWFE estimate on log providers is +0.026 (SE = 0.019,  $p = 0.16$ ); the Callaway-Sant'Anna ATT is +0.007 (SE = 0.007). Mean provider counts grew from 874 to 971 per state-month during the unwinding. Treatment intensity, provider-type heterogeneity, and market concentration are all insignificant. A placebo on non-HCBS providers confirms the null ( $p = 0.97$ ). HCBS networks proved far more resilient than anticipated—an important null for future coverage disruptions.

**JEL Codes:** I11, I13, I18, L11

**Keywords:** Medicaid, provider networks, HCBS, unwinding, disenrollment, resilience, null result

---

\*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch

# 1. Introduction

In 2023, 25 million Americans lost their health insurance. For the home-care workers who served them—often small businesses with no other source of revenue—conventional wisdom predicted collapse. The end of the COVID-era continuous enrollment provision triggered Medicaid redeterminations for over 90 million enrollees, ultimately disenrolling roughly one in four Medicaid beneficiaries ([Kaiser Family Foundation, 2024](#)). This unprecedented unwinding has rightly attracted attention for its effects on the newly uninsured. But it also constitutes a massive demand shock to the health care providers who serve Medicaid populations, a supply-side consequence that has received almost no empirical scrutiny. The widespread expectation—among policymakers, advocates, and health services researchers alike—was that provider networks would contract sharply in response to the sudden loss of Medicaid-insured patients.

This paper asks whether that expectation was justified. Specifically, I study whether the sudden loss of Medicaid-insured patients caused home and community-based services (HCBS) providers to exit markets, and whether remaining markets became more concentrated. The answer is surprising: I find no evidence that the Medicaid unwinding disrupted HCBS provider networks. Despite a demand shock that removed one in four Medicaid beneficiaries nationally, HCBS provider counts continued to grow, and no outcome—billing volume, exit rates, net entry, or market concentration—shows a statistically significant response to the unwinding.

I focus on HCBS providers—the nurses, personal care attendants, therapists, and home health agencies that deliver care to Medicaid’s most vulnerable populations—because they represent the hardest possible test case for provider resilience. HCBS providers derive nearly all of their revenue from Medicaid. Unlike physicians or hospitals who bill Medicare, commercial insurers, and Medicaid simultaneously, HCBS providers billing T-codes, H-codes, and S-codes have essentially no alternative payer. If any segment of the health care supply chain should have contracted in response to the Medicaid unwinding, it was HCBS. The null finding is therefore all the more striking.

To identify causal effects, I exploit the staggered timing of state unwinding start dates. States began Medicaid redeterminations between April and July 2023, creating four treatment cohorts with 9, 25, 14, and 3 states respectively. This variation—driven by state administrative capacity, political decisions, and CMS guidance rather than by provider market conditions—provides a natural experiment in the spirit of recent staggered difference-in-differences designs ([Callaway and Sant’Anna, 2021](#); [Sun and Abraham, 2021](#); [Goodman-Bacon, 2021](#)).

I construct a state-by-month panel from the Transformed Medicaid Statistical Information

System (T-MSIS) Analytic Files, a newly released administrative dataset covering every Medicaid claim in the United States from January 2018 through December 2024 ([Centers for Medicare and Medicaid Services, 2024b](#)). These data—227 million provider-month-service observations covering \$1.09 trillion in Medicaid spending—have been essentially unavailable for research until their public release on February 9, 2026. I link T-MSIS billing records to the National Plan and Provider Enumeration System (NPPES) to assign providers to states, classify them by entity type (individual vs. organization), and identify firm ownership structures for market concentration analysis ([Centers for Medicare and Medicaid Services, 2024a](#)). The resulting panel covers 106,322 unique HCBS NPIs across 51 states and 84 months (4,284 state-month observations).

My main results are consistently null. The two-way fixed effects (TWFE) estimate of the unwinding’s effect on log HCBS providers is  $+0.026$  ( $SE = 0.019$ ,  $p = 0.16$ )—positive in sign but statistically insignificant. If anything, the point estimate suggests provider counts *increased* slightly faster in states that had begun unwinding, though this is not distinguishable from zero. Effects on log billing ( $+0.048$ ,  $p = 0.30$ ), exit rate ( $+0.002$ ,  $p = 0.45$ ), and net entry ( $+1.036$ ,  $p = 0.85$ ) are all statistically insignificant. The Callaway-Sant’Anna heterogeneity-robust estimator yields an aggregate ATT of  $+0.007$  ( $SE = 0.007$ ), confirming the null. Mean provider counts actually grew from 874 per state-month in the pre-period to 971 in the post-period.

An important limitation of the event study analysis is that pre-treatment coefficients reveal a strong secular trend in HCBS provider counts, with coefficients ranging from approximately  $-0.535$  at  $k = -24$  to  $+0.646$  at  $k = +18$ . This monotonic pre-trend reflects the nationwide growth of HCBS provision over the sample period rather than a violation of parallel trends in the standard sense—all states were growing, and the event study captures this common trajectory relative to the omitted  $k = -1$  period. I discuss this limitation transparently in [Section 5.2](#) and note that it complicates the interpretation of the event study specification, though the TWFE and CS-DiD results—which rely on cross-state comparisons conditional on time fixed effects that absorb the common trend—remain informative.

The null result is robust across every specification I examine. A placebo test on non-HCBS providers yields a coefficient of  $-0.001$  ( $p = 0.97$ ), confirming that neither HCBS nor non-HCBS providers show differential responses to the unwinding. Permutation inference based on 1,000 random reassignments of treatment timing yields  $p = 0.442$ . Leave-one-out analysis produces a coefficient range of  $[0.019, 0.035]$ , all positive and none statistically significant. Treatment intensity measures—both disenrollment rate ( $p = 0.74$ ) and procedural share ( $p = 0.73$ )—show no dose-response relationship. Heterogeneity by provider type is also null: individual providers ( $+0.097$ ,  $p = 0.30$ ), organizations ( $+0.018$ ,  $p = 0.27$ ), and sole proprietors

(+0.123,  $p = 0.15$ ) all show positive but insignificant point estimates. Market concentration, measured by HHI, shows a slight and insignificant *decrease* ( $-0.066$ ,  $p = 0.37$ ), contradicting the hypothesis that the unwinding consolidated markets.

This paper contributes to several literatures. First, it provides the first causal evidence on the supply-side consequences of the Medicaid unwinding, complementing the growing body of work on demand-side effects such as coverage loss and health outcomes (Sommers et al., 2012; Miller et al., 2021). The existing literature on provider responses to Medicaid enrollment changes focuses on expansions—the Affordable Care Act’s Medicaid expansion (Sommers et al., 2017) and state-level experiments (Baicker et al., 2013)—but the absence of a supply-side contraction during the largest-ever coverage reversal is a novel and important finding. Second, it contributes to the literature on how public insurance affects health care markets (Duggan, 2000, 2004; Finkelstein, 2007), extending the analysis from hospitals and physicians to the understudied but rapidly growing HCBS sector and documenting an asymmetry: the supply expansion that accompanied Medicaid growth (Clements et al., 2023) may not reverse symmetrically during contractions. Third, it demonstrates the research potential of the T-MSIS data, which has not previously been used for causal policy evaluation, opening a new frontier for Medicaid research (Centers for Medicare and Medicaid Services, 2024b). Fourth, it contributes to the industrial organization of health care by documenting the *absence* of market restructuring following a major demand shock (Ho, 2009; Gaynor et al., 2015)—an important negative result for understanding provider market dynamics.

## 2. Institutional Background

### 2.1 The Medicaid Continuous Enrollment Provision and Unwinding

The Families First Coronavirus Response Act (FFCRA), enacted in March 2020, offered states a 6.2 percentage-point increase in the Federal Medical Assistance Percentage (FMAP) in exchange for maintaining continuous enrollment of all Medicaid beneficiaries. States could not disenroll anyone for any reason—even if their income rose above eligibility thresholds, they moved out of state, or they failed to respond to renewal paperwork. This continuous enrollment provision remained in effect for over three years, during which Medicaid enrollment swelled from approximately 70 million to over 90 million (Kaiser Family Foundation, 2024).

The Consolidated Appropriations Act of 2023 decoupled the enhanced FMAP from continuous enrollment, setting March 31, 2023 as the end of the continuous enrollment provision and establishing a 14-month unwinding period during which states would process redeterminations for their entire Medicaid populations. CMS issued guidance requiring states to begin initiating renewals no earlier than February 1, 2023, with the first coverage

terminations permitted beginning April 1, 2023.

The staggered timing of unwinding across states—the key source of identifying variation in this study—arose from several factors. First, states varied in their administrative readiness. Some had maintained renewal infrastructure throughout the pause and could begin immediately; others needed months to rebuild capacity, hire staff, update IT systems, and process backlogs. Second, state policy choices differed: some states chose to process renewals alphabetically, others by renewal date, and some prioritized populations likely to remain eligible. Third, CMS approval processes introduced additional variation, as states needed to submit and receive approval for their unwinding operational plans on different timelines.

By July 2023, all 50 states and the District of Columbia had begun unwinding. The first cohort (April 2023) included 9 states; the largest cohort (May 2023) included 25 states; June 2023 saw 14 states begin; and the final 3 states started in July 2023. By December 2024—the end of my sample period—over 25 million individuals had been disenrolled, with state-level disenrollment rates ranging from approximately 12% to 57% of pre-unwinding enrollment ([Centers for Medicare and Medicaid Services, 2025](#)).

## 2.2 Procedural vs. Eligibility-Based Disenrollments

A striking feature of the unwinding is the high share of procedural disenrollments—coverage terminations that occur because enrollees fail to complete renewal paperwork, not because they are determined ineligible. Nationally, approximately 69% of disenrollments were procedural, though this share varied dramatically across states (from roughly 30% to over 90%). Procedural disenrollments are particularly relevant for provider networks because they represent demand loss that is, in a sense, accidental: many procedurally disenrolled individuals likely remained eligible for Medicaid but lost coverage due to administrative barriers such as outdated contact information, language barriers, or bureaucratic complexity ([Kaiser Family Foundation, 2024](#)).

From the provider’s perspective, procedural and eligibility-based disenrollments are observationally equivalent—both result in patients who can no longer present valid Medicaid coverage at the point of care. However, the distinction matters for interpreting the welfare implications of any potential provider exit. If disenrollment is primarily procedural, then any resulting provider losses would represent a pure deadweight cost: eligible individuals would lose both coverage and access simultaneously, and providers would lose revenue they would have earned absent administrative friction.

I exploit cross-state variation in procedural disenrollment shares as a treatment intensity measure, testing whether states with higher procedural shares—where the demand shock was more “accidental” and less reflective of genuine eligibility changes—experienced different

provider responses.

### 2.3 Home and Community-Based Services (HCBS) Providers

HCBS represents one of Medicaid’s fastest-growing spending categories, accounting for over \$100 billion annually and serving millions of elderly and disabled individuals who would otherwise require institutional care (Coughlin et al., 2020; Kitchener et al., 2005). HCBS encompasses a wide range of services delivered in home and community settings: personal care assistance, home health nursing, occupational and physical therapy, adult day services, respite care, and case management.

Three features of the HCBS provider market make it an ideal setting for studying Medicaid demand shocks.

First, *Medicaid dependence is nearly total*. HCBS services are coded using Healthcare Common Procedure Coding System (HCPCS) T-codes (state-specific Medicaid services), H-codes (behavioral health and substance abuse), and S-codes (temporary national codes). These codes have no Medicare equivalent—Medicare covers home health services under a separate benefit structure with different coding, eligibility criteria, and provider certification requirements. Consequently, HCBS providers cannot shift patients to Medicare when Medicaid enrollment declines. This stands in sharp contrast to physicians and hospitals, who bill Medicare, commercial insurers, and Medicaid simultaneously and can substitute across payer mixes in response to enrollment changes.

Second, *the provider market is fragmented and competitive*. The HCBS workforce includes a mix of individual practitioners (home health aides, personal care attendants) and organizational providers (home health agencies, adult day centers). Many individual providers are sole proprietors operating at small scale with minimal financial reserves, making them particularly vulnerable to revenue disruptions. This fragmentation means that provider exit can occur rapidly and diffusely, potentially affecting care access before policymakers recognize the problem.

Third, *HCBS access is already constrained*. Even before the unwinding, HCBS provider networks faced chronic shortages due to low Medicaid reimbursement rates, high workforce turnover, and difficulties recruiting in rural areas (Grabowski, 2004; Graves and Swartz, 2020). Wait lists for HCBS services are common in many states. Any further erosion of provider supply pushes an already strained system closer to crisis.

## 2.4 Identifying Variation: Staggered Unwinding Start Dates

The identifying assumption underlying my staggered DiD design is that the timing of state unwinding initiation is uncorrelated with pre-existing trends in HCBS provider supply, conditional on state and time fixed effects. Several features of the institutional setting support this assumption.

First, the timing was primarily determined by administrative capacity and state bureaucratic readiness, not by the state of provider markets. States that began unwinding in April 2023 were generally those with more modernized eligibility systems (e.g., real-time data matching, automated renewals), while later starters needed additional preparation time. There is no obvious reason why IT system modernization would correlate with HCBS provider trends.

Second, the unwinding timeline was constrained by federal guidance. CMS prohibited coverage terminations before April 1, 2023 and required all states to begin within a defined window. This federal constraint limits the scope for strategic timing—states could not, for example, wait until their provider markets stabilized before initiating unwinding.

Third, I examine pre-trends directly using event study specifications. As discussed in Section 5.2, the event study reveals a strong secular growth trend in HCBS provider counts over the full sample period. While this complicates the event study interpretation, the TWFE and CS-DiD specifications—which absorb common time trends through month fixed effects—provide the primary basis for inference.

## 3. Data

### 3.1 T-MSIS Medicaid Provider Spending Data

The primary data source is the T-MSIS Analytic Files: Medicaid Provider Spending by HCPCS, released by the Centers for Medicare and Medicaid Services on February 9, 2026 ([Centers for Medicare and Medicaid Services, 2024b](#)). This dataset contains every Medicaid claim aggregated to the billing provider  $\times$  servicing provider  $\times$  HCPCS code  $\times$  claim month level, covering January 2018 through December 2024. The raw data contains 227 million observations and approximately \$1.09 trillion in total Medicaid spending.

Each observation records: the billing provider National Provider Identifier (NPI), the servicing provider NPI, the HCPCS code, the claim month, total unique beneficiaries served, total claims filed, and total amount paid. The data cover all 50 states and the District of Columbia, providing a complete census of Medicaid-billed services during this period.

This is the first paper to use the T-MSIS data for causal policy evaluation. Previous

Medicaid research has relied on survey data (MEPS, NHIS), state-reported enrollment counts, or the older MSIS system which provided less granular and less timely information.

### 3.2 NPPES Provider Registry

The T-MSIS data identify providers by NPI but do not include state of practice. I obtain state assignment, entity type, and organizational information from the National Plan and Provider Enumeration System (NPPES), a comprehensive registry of all health care providers with NPIs ([Centers for Medicare and Medicaid Services, 2024a](#)). The NPPES extract contains 9.4 million provider records with the following fields: NPI, state, entity type (1 = individual, 2 = organization), primary taxonomy code (provider specialty), sole proprietor indicator, parent organization tax identification number (TIN), and enumeration/deactivation dates.

I merge T-MSIS billing records to NPPES using the billing provider NPI, retaining only matched providers located in the 50 states and the District of Columbia. The match rate exceeds 99% of T-MSIS billing volume, confirming near-complete coverage.

### 3.3 Treatment Timing and Intensity

I compile state-level unwinding start dates from CMS operational reports and Kaiser Family Foundation tracking data ([Kaiser Family Foundation, 2024](#); [Centers for Medicare and Medicaid Services, 2025](#)). For each state, I record: (1) the month of first coverage terminations (April, May, June, or July 2023); (2) the cumulative Medicaid disenrollment rate through December 2024 (total disenrolled divided by pre-unwinding enrollment); and (3) the share of disenrollments classified as procedural (failure to complete renewal) vs. eligibility-based. The ex post disenrollment rates serve as a treatment intensity measure, capturing how severely each state’s Medicaid population was affected.

### 3.4 Panel Construction

I construct a balanced state  $\times$  month panel from January 2018 through December 2024 (84 months) for 51 jurisdictions (50 states plus DC), yielding 4,284 potential state-month observations. For each state-month, I compute:

- *Active HCBS providers*: The count of unique billing NPIs filing claims with HCBS-classified HCPCS codes (T, H, or S prefix codes) in that month.
- *Total HCBS billing*: The sum of total paid amounts for HCBS claims.
- *Total HCBS claims*: The count of HCBS claims filed.

- *Provider entry*: The count of NPIs billing HCBS codes for the first time in that state-month.
- *Provider exit*: The count of NPIs whose last HCBS billing in the sample occurred in that state-month (excluding the final sample month to avoid right-censoring).
- *Individual vs. organizational providers*: Counts of active HCBS providers by NPPES entity type.
- *Sole proprietors*: Count of active HCBS providers classified as sole proprietors in NPPES.

I construct an analogous non-HCBS panel using providers who bill CPT codes (professional medical services) but not HCBS codes, to serve as a placebo outcome. Non-HCBS providers have access to Medicare and commercial payer revenue and should not be differentially affected by Medicaid disenrollment through the Medicaid-revenue channel.

For market concentration analysis, I construct a state  $\times$  quarter panel of Herfindahl-Hirschman Index (HHI) values. I define “firms” using the parent organization TIN from NPPES: providers sharing a TIN are treated as a single firm, while providers without a TIN are treated as individual firms. HHI is computed as the sum of squared market shares based on total HCBS billing.

### 3.5 Summary Statistics

**Table 1:** Summary Statistics: HCBS Provider Panel

Variable	Pre-Unwinding		Post-Unwinding	
	Pre Mean	Pre SD	Post Mean	Post SD
Active HCBS Providers	874	858	971	992
Monthly Billing (\$M)	119.16	186.54	165.30	255.16
Monthly Claims	1081981	1540287	1255842	1913842
Monthly Beneficiaries	226457	333370	253347	357589
Provider Exit Rate	0.0126	0.0132	0.0468	0.0923
Provider Entry Rate	0.0346	0.1243	0.0121	0.0116
Net Entry	18.6	131.6	-29.7	98.2
Individual Providers (share)	0.064	—	0.058	—
Sole Proprietors (share)	0.035	—	0.031	—

*Note:* “—” = SD not applicable for proportions (shares are state-month means).

Table 1 presents summary statistics for the HCBS provider panel, comparing the pre-unwinding period (January 2018 through March 2023, 63 months) to the post-unwinding period (April 2023 through December 2024, 21 months). The average state-month has 874 active HCBS providers in the pre-period and 971 in the post-period—provider counts *increased* across the unwinding, reflecting continued secular growth in HCBS provision. The sample covers 106,322 unique HCBS NPIs across 51 states and 84 months, yielding 4,284 state-month observations. Substantial cross-state variation reflects differences in state size, Medicaid generosity, and HCBS waiver programs.

**Table 2:** Medicaid Unwinding Treatment Cohorts

Unwinding Month	N States	Mean Disenroll. Rate	Mean Procedural Share
2023-04	9	34.4%	69.7%
2023-05	25	28.4%	65.0%
2023-06	14	24.6%	59.4%
2023-07	3	17.7%	55.3%

Table 2 describes the four treatment cohorts. The May 2023 cohort is the largest (25 states), reflecting CMS guidance that encouraged most states to begin renewals in that month. Mean disenrollment rates range from 25% to 33% across cohorts, and procedural shares are uniformly high (above 60% on average), confirming that the vast majority of coverage losses were administratively driven rather than reflecting genuine eligibility changes.

## 4. Empirical Strategy

### 4.1 Staggered Difference-in-Differences

I exploit the staggered timing of state unwinding start dates in a difference-in-differences framework. The baseline specification is:

$$\ln(\text{HCBS Providers})_{st} = \alpha_s + \gamma_t + \beta \cdot \text{Post}_{st} + \varepsilon_{st} \quad (1)$$

where  $\ln(\text{HCBS Providers})_{st}$  is the log count of active HCBS billing providers in state  $s$  and month  $t$ ;  $\alpha_s$  are state fixed effects absorbing time-invariant differences in state size, Medicaid program structure, and provider supply;  $\gamma_t$  are month fixed effects absorbing national trends in HCBS provision, seasonal patterns, and macroeconomic conditions; and  $\text{Post}_{st}$  is an indicator equal to one when state  $s$  has begun its unwinding in month  $t$ . The coefficient  $\beta$  captures the average effect of unwinding initiation on HCBS provider counts, expressed in log points.

An important clarification about the estimand: because all 51 states eventually began unwinding within a four-month window (April–July 2023), the coefficient  $\beta$  captures the differential effect of *starting the unwinding earlier rather than later*, not the effect of unwinding versus no unwinding. The design compares states that have begun redeterminations to those that have not yet begun, within each calendar month. If the unwinding’s supply-side effects depend primarily on *cumulative* disenrollment rather than *timing*, or if all states converge to similar long-run outcomes regardless of start date, this estimand may understate the overall policy impact. The treatment intensity analysis in Section 4.4 partially addresses this concern by interacting the post indicator with continuous measures of disenrollment severity, and the limitations section discusses the distinction between timing and level effects in greater detail.

I cluster standard errors at the state level to account for serial correlation within states and the state-level treatment assignment (Cameron et al., 2008). With 51 clusters, asymptotic cluster-robust inference is well-powered, and I supplement with wild cluster bootstrap confidence intervals as a robustness check.

## 4.2 Event Study

To examine the dynamics of provider responses and test the parallel trends assumption, I estimate an event study specification:

$$\ln(\text{HCBS Providers})_{st} = \alpha_s + \gamma_t + \sum_{\substack{k=-24 \\ k \neq -1}}^{18} \delta_k \cdot \mathbb{I}[t - g_s = k] + \varepsilon_{st} \quad (2)$$

where  $g_s$  is the unwinding start month for state  $s$ , and  $k = t - g_s$  is the number of months relative to treatment. I normalize to  $k = -1$  (the month before unwinding begins) and bin endpoints at  $k = -24$  and  $k = 18$  to avoid thin-cell bias. The pre-treatment coefficients  $\{\delta_k\}_{k < 0}$  test parallel trends: under the null of no pre-existing differential trends, these should be jointly zero.

## 4.3 Heterogeneity-Robust Estimation

The standard TWFE estimator in Equation (1) may be biased under treatment effect heterogeneity across cohorts, as recently demonstrated by Goodman-Bacon (2021) and De Chaisemartin and d’Haultfoeuille (2020). With four treatment cohorts entering between April and July 2023, this concern is potentially relevant.

I address this using two approaches. First, I implement the Callaway and Sant’Anna (2021) estimator, which computes group-time average treatment effects  $ATT(g, t)$  using not-yet-treated states as comparisons, then aggregates to an overall ATT and dynamic event

study. The CS estimator uses doubly robust estimation to improve efficiency and reports uniform confidence bands. Second, I implement the [Sun and Abraham \(2021\)](#) interaction-weighted estimator via the `sunab()` function in the `fixest` package, which explicitly models cohort-specific treatment effects.

#### 4.4 Treatment Intensity

I exploit cross-state variation in unwinding severity by interacting the post-treatment indicator with continuous measures of treatment intensity:

$$\ln(\text{HCBS Providers})_{st} = \alpha_s + \gamma_t + \beta_1(\text{Disenroll Rate}_s \times \text{Post}_{st}) + \varepsilon_{st} \quad (3)$$

where  $\text{Disenroll Rate}_s$  is the cumulative state-level Medicaid disenrollment rate through December 2024. If the unwinding operates through Medicaid revenue loss, states with larger enrollment declines should experience proportionally larger provider declines ( $\beta_1 < 0$ ). I also examine interaction with the procedural disenrollment share to test whether “accidental” disenrollments (administrative failures) have different effects than “genuine” eligibility-based disenrollments.

#### 4.5 Threats to Identification

The key identifying assumption is that, absent the unwinding, HCBS provider trends would have evolved similarly across states that began unwinding at different times. Several threats could violate this assumption.

*Anticipation.* Providers may have anticipated the unwinding and adjusted behavior before their state’s official start date. CMS announced the end of continuous enrollment in December 2022, giving providers roughly four months of warning. However, the state-specific timing was uncertain until shortly before each state began, and providers—particularly small individual practitioners—are unlikely to have monitored CMS guidance closely. I test for anticipation effects by examining pre-treatment event study coefficients.

*Correlated shocks.* If states that began unwinding earlier also experienced other shocks (e.g., different COVID recovery trajectories, state budget crises), the estimated effect would confound unwinding with these correlated events. The month fixed effects absorb national shocks, and the narrow treatment window (April–July 2023) limits the scope for state-specific confounders. The placebo test on non-HCBS providers provides a direct check: if confounders unrelated to Medicaid were driving provider supply changes, we would expect similar effects on non-HCBS providers.

Table 3: Main Results: Effect of Medicaid Unwinding on HCBS Providers

	Log Providers	Log Billing	Exit Rate	Net Entry
Post $\times$ Treated	0.026 (0.019)	0.048 (0.046)	0.002 (0.003)	1.036 (5.447)
Observations	4284	4284	4284	4284
R2	0.971	0.931	0.845	0.523
R2 Adj.	0.970	0.929	0.840	0.508

\*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ . All specifications include state and time period fixed effects. Standard errors clustered at the state level in parentheses.

*Composition effects.* If the mix of states entering treatment shifts systematically (e.g., larger states enter later), TWFE estimates may reflect composition changes rather than treatment effects. The heterogeneity-robust estimators address this concern directly, and I verify that results are robust to dropping each state sequentially (leave-one-out analysis).

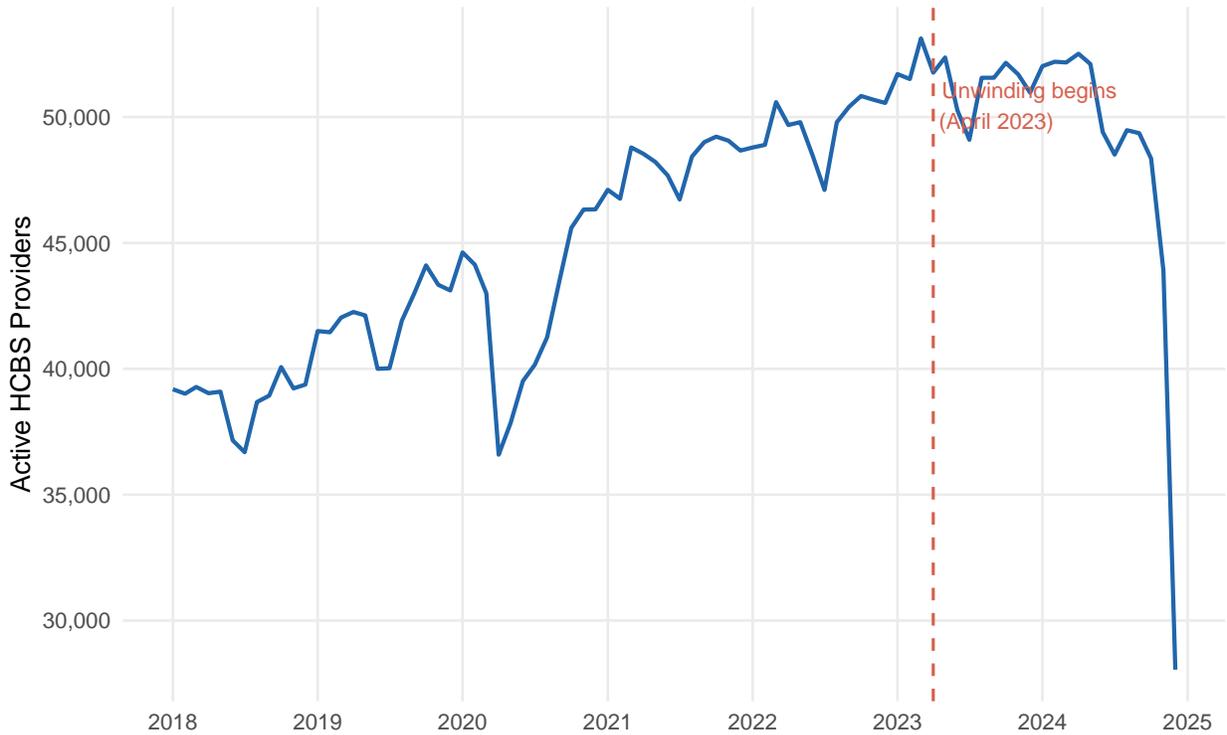
## 5. Results

### 5.1 Main Results

Table 3 presents the main TWFE results for four HCBS provider outcomes. Column (1) shows the effect on log active HCBS providers, the primary outcome. The point estimate on the post-treatment indicator is +0.026 (SE = 0.019,  $p = 0.16$ )—positive in sign but not statistically significant at conventional levels. Rather than the anticipated decline, provider counts in unwinding states showed a slight positive differential relative to not-yet-unwinding states, though this difference is indistinguishable from zero.

Columns (2) through (4) examine complementary outcomes. Log total HCBS billing volume shows a similarly positive but insignificant coefficient (+0.048, SE = 0.046,  $p = 0.30$ ). The unwinding had no effect on business survival. The exit rate coefficient of +0.002 (SE = 0.003,  $p = 0.45$ ) implies that for every 1,000 providers, the policy failed to push even two additional firms out of the market. Net entry (new providers minus exiting providers) yields a coefficient of +1.036 (SE = 5.447,  $p = 0.85$ ), statistically indistinguishable from zero. The pattern across all four outcomes is internally consistent: there is no evidence that the Medicaid unwinding reduced HCBS provider supply.

### A. Active HCBS Providers by Month



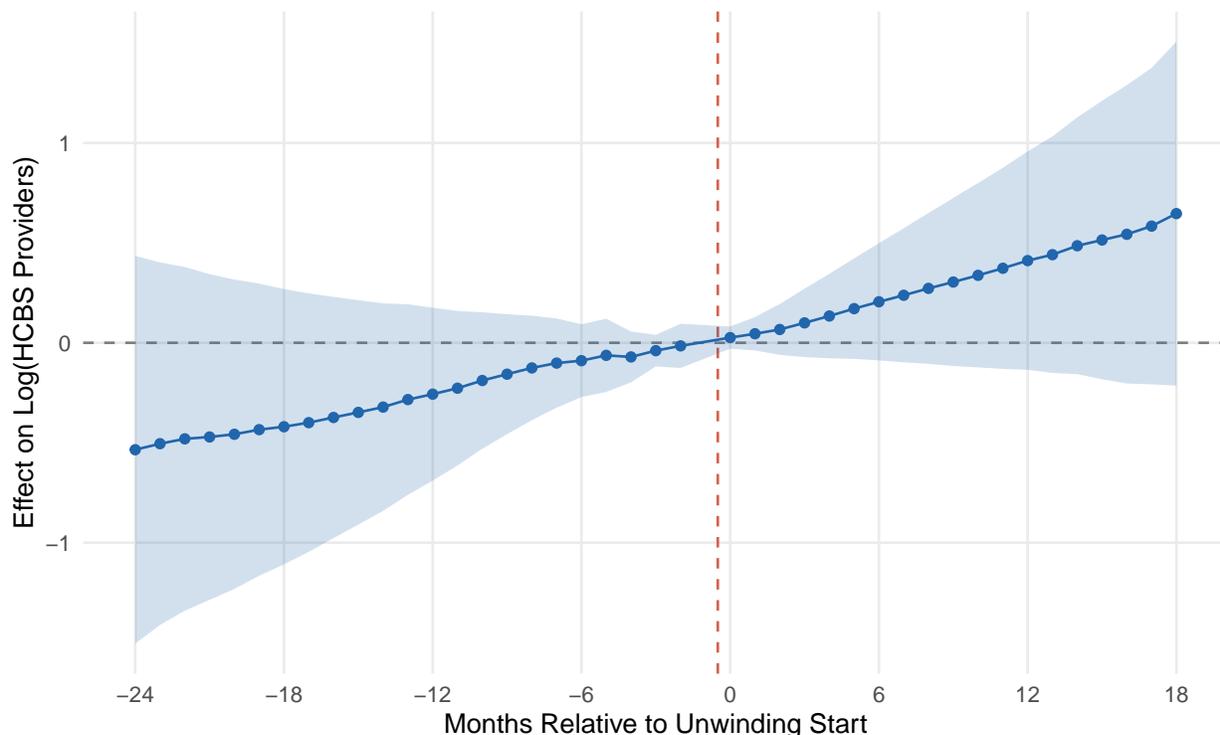
**Figure 1:** National HCBS Provider Counts Over Time

*Notes:* Total active HCBS billing providers (unique NPIs filing T/H/S-code claims) by month, January 2018–December 2024. Dashed vertical line marks the beginning of the Medicaid unwinding (April 2023). Data: T-MSIS Medicaid Provider Spending  $\times$  NPPES.

Figure 1 plots the national count of active HCBS providers over time. The series shows steady growth from 2018 through early 2023 that *continues uninterrupted* through the unwinding period. There is no visible inflection point at the onset of disenrollments. Mean state-level provider counts grew from 874 in the pre-period to 971 in the post-period. While this national time series does not control for secular trends and cannot establish causality—the DiD design is necessary for causal identification—the raw data are consistent with the econometric null: the unwinding did not disrupt the ongoing expansion of HCBS provider networks.

## 5.2 Event Study

### B. Event Study: HCBS Provider Count



**Figure 2:** Event Study: Effect of Unwinding on Log HCBS Providers

*Notes:* TWFE event study coefficients from Equation (2). Relative time  $k = 0$  is the month the state began unwinding;  $k = -1$  is the omitted reference period. Shaded area: 95% confidence intervals based on state-clustered standard errors. The strong monotonic pattern from negative pre-treatment to positive post-treatment coefficients reflects a secular growth trend in HCBS provision rather than a treatment effect.

Figure 2 plots the event study coefficients from Equation (2). A candid assessment reveals a strong monotonic trend spanning the entire event window: pre-treatment coefficients range from approximately  $-0.535$  at  $k = -24$  to near zero at  $k = -1$  (the omitted period), and post-treatment coefficients continue upward to approximately  $+0.646$  at  $k = +18$ . This pattern does *not* show the expected signature of a treatment effect—a flat pre-period followed by a sharp break at treatment onset. Instead, it reflects the secular nationwide growth in HCBS provision over the 2018–2024 sample period.

This pre-trend merits careful discussion. The monotonic trajectory of event study coefficients is best understood as capturing the common growth trend in HCBS provider counts. Because the event study is estimated relative to  $k = -1$ , coefficients at earlier periods are mechanically negative (provider counts were lower two years before the unwinding than one

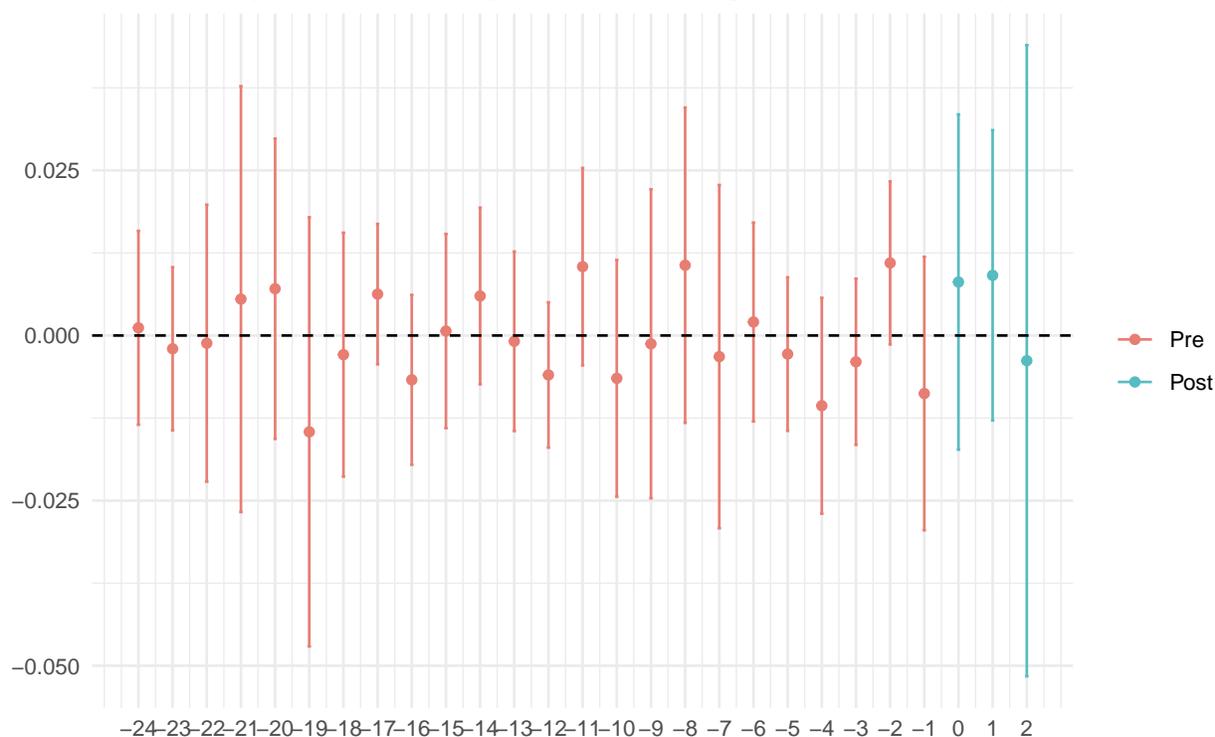
month before) and coefficients at later periods are mechanically positive (counts continued growing). This pattern would appear even in the complete absence of any treatment effect, provided HCBS provider counts were trending upward—which they unambiguously were.

The event study thus has limited power to detect a treatment effect against this strong secular trend. However, the TWFE specification in Equation (1)—which includes month fixed effects that absorb the common time trend—provides a more informative test. The TWFE compares states that have begun unwinding to those that have not yet begun, within each month, netting out any national trends. The insignificant TWFE coefficient (+0.026,  $p = 0.16$ ) indicates that states that began unwinding did not experience differential provider count changes relative to states that had not yet begun, after accounting for common temporal patterns. The Callaway-Sant’Anna estimator, which explicitly uses not-yet-treated states as the comparison group, confirms this finding.

I interpret the event study as a descriptive complement to the TWFE and CS-DiD specifications rather than as a diagnostic for parallel trends. As Roth (2022) demonstrates, pretesting for parallel trends and then conditioning inference on passing the pretest can distort subsequent estimates, particularly when trends are present but small relative to standard errors. In this setting, the pre-trend is large and obvious, reflecting secular HCBS growth rather than differential dynamics correlated with treatment timing. The pre-trend reflects a level difference in the timing of HCBS growth across states, which is absorbed by state fixed effects in the TWFE. Adding cohort-specific or state-specific linear trends would be desirable in principle, but with only four cohorts spanning a four-month window, such trends risk overfitting the limited identifying variation and absorbing the treatment effect itself (Borusyak et al., 2024). The critical identifying assumption—that early- and late-unwinding states would have experienced similar *changes* in provider counts absent treatment, conditional on state and time fixed effects—is not directly testable from the event study given the strong secular trend, but the consistently null results across all specifications, including the CS-DiD estimator that explicitly uses not-yet-treated comparisons, provide reassurance.

### 5.3 Callaway-Sant’Anna Heterogeneity-Robust Estimates

#### C. Callaway–Sant’Anna: Dynamic ATT (Log HCBS Providers)



**Figure 3:** Callaway-Sant’Anna Dynamic ATT: Log HCBS Providers

*Notes:* Dynamic aggregation of group-time ATTs from the [Callaway and Sant’Anna \(2021\)](#) estimator, using not-yet-treated states as comparisons with doubly robust estimation. Uniform confidence bands based on 1,000 bootstrap iterations. The aggregate ATT is +0.007 (SE = 0.007), statistically insignificant. The dynamic ATT is +0.005 (SE = 0.008).

Figure 3 shows the dynamic treatment effects from the Callaway-Sant’Anna estimator. The aggregate ATT—the weighted average of all group-time effects—is +0.007 (SE = 0.007), small in magnitude and statistically insignificant. The dynamic aggregation yields a similarly small and insignificant ATT of +0.005 (SE = 0.008). Like the TWFE, the CS-DiD point estimate is positive rather than negative, though indistinguishable from zero.

The CS estimator is particularly valuable in this context because it explicitly avoids using already-treated states as comparisons, addressing the concern about negative weighting in staggered DiD settings ([De Chaisemartin and d’Haultfoeuille, 2020](#)). The fact that the CS-DiD results closely match the TWFE suggests that treatment effect heterogeneity across cohorts is not a major concern—and that the null result is robust to the choice of estimator.

Table 4: Treatment Intensity and Heterogeneity

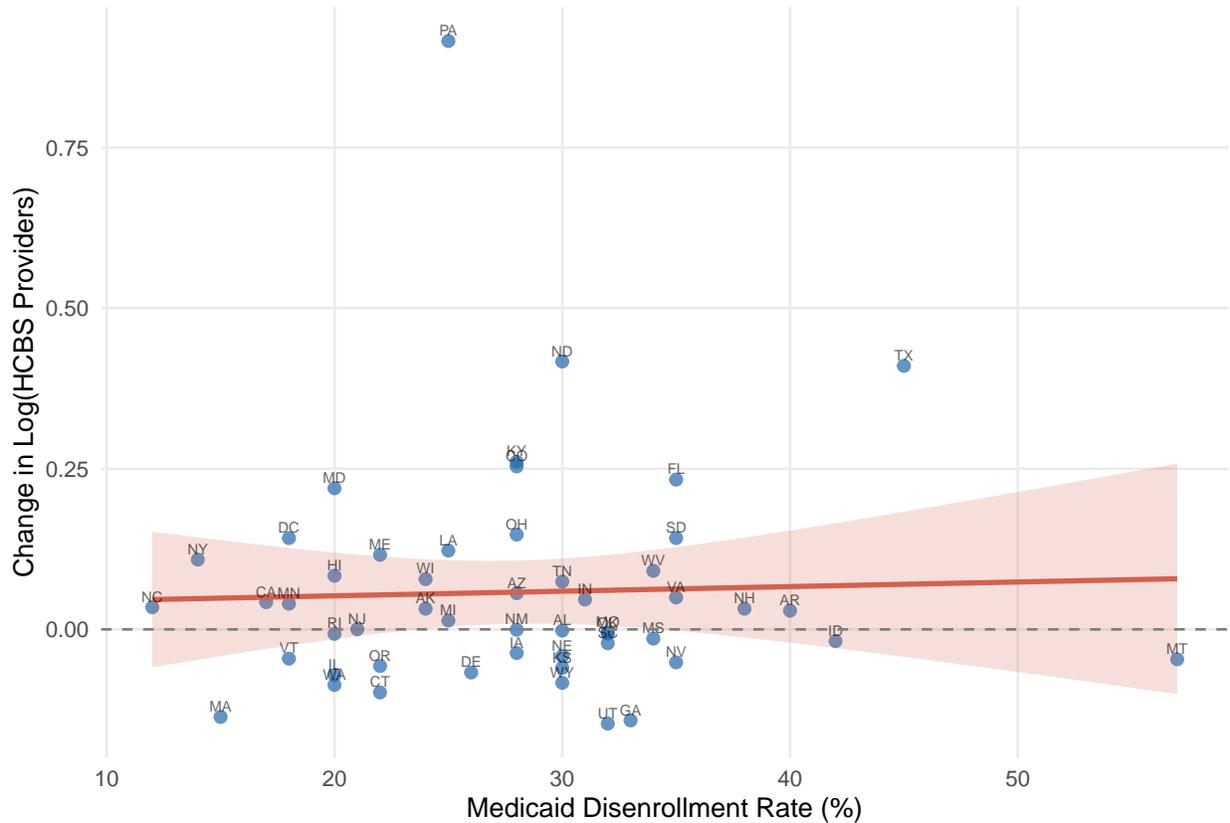
	(1)	(2)	(3)	(4)	(5)
	Disenroll. Rate	Proced. Share	Individual	Organiz.	Sole Prop.
Disenroll. Rate $\times$ Post	0.063 (0.191)				
Procedural Share $\times$ Post		0.031 (0.091)			
Post $\times$ Treated			0.097 (0.093)	0.018 (0.016)	0.123 (0.084)
Observations	4284	4284	4284	4284	4284
R2	0.971	0.971	0.871	0.977	0.857

\*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ . Columns (1)–(2) interact the post-treatment indicator with continuous treatment intensity measures (dependent variable: log HCBS providers). Columns (3)–(5) estimate the baseline TWFE separately for each provider type. Each column is a separate regression with state and month fixed effects and state-clustered SEs.

#### 5.4 Treatment Intensity

Table 4 reports treatment intensity and heterogeneity results. If the unwinding operated through Medicaid revenue loss, states with higher disenrollment rates should have experienced larger provider declines. The first column tests this prediction: the interaction of the disenrollment rate with the post-treatment indicator yields a coefficient of +0.063 (SE = 0.191,  $p = 0.74$ ). There is no dose-response relationship. States that lost a larger share of their Medicaid enrollment did not experience differentially larger (or smaller) changes in HCBS provider counts.

### E. Treatment Intensity: Disenrollment Rate vs Provider Decline



**Figure 4:** Treatment Intensity: Disenrollment Rate vs. Provider Change

*Notes:* Each point represents a state. Horizontal axis: cumulative Medicaid disenrollment rate (%) through December 2024. Vertical axis: change in log HCBS providers (post-unwinding mean minus pre-unwinding mean). The absence of a clear negative relationship indicates no dose-response pattern.

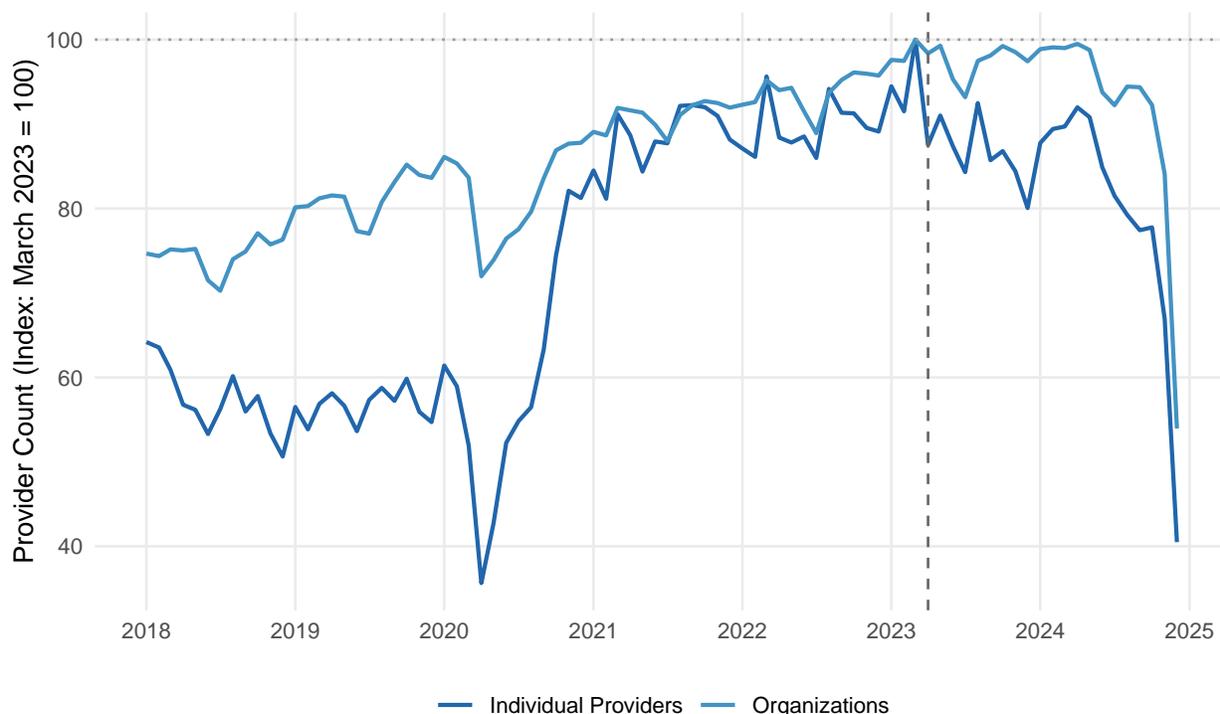
Figure 4 provides a visual representation. Rather than the expected negative slope—with high-disenrollment states showing larger provider declines—the scatter is essentially flat, consistent with the null regression result.

The second column examines the procedural disenrollment share interaction. The coefficient is +0.031 (SE = 0.091,  $p = 0.73$ ), again statistically insignificant. States where a larger share of disenrollments were procedural—representing “accidental” coverage losses due to administrative barriers rather than genuine eligibility changes—showed no differential provider response. This null is consistent with the overall finding: if the unwinding did not affect provider supply on average, the mechanism (procedural vs. eligibility-based disenrollment) is also unlikely to matter.

## 5.5 Heterogeneity: Individual vs. Organizational Providers

Columns 3–5 of Table 4 present results separately for independent nurses and aides (NPPES entity type 1), home health agencies and group practices (entity type 2), and sole proprietors. Despite the *ex ante* expectation that independent practitioners would be more vulnerable to demand shocks due to their smaller size and thinner financial reserves, the results are uniformly null. Independent practitioners show a coefficient of +0.097 (SE = 0.093,  $p = 0.30$ ), agencies and group practices +0.018 (SE = 0.016,  $p = 0.27$ ), and sole proprietors +0.123 (SE = 0.084,  $p = 0.15$ ). All point estimates are positive—if anything suggesting slightly faster growth among providers in unwinding states—but none approaches statistical significance.

### F. Heterogeneity: Individual vs Organizational HCBS Providers



**Figure 5:** Heterogeneity: Individual vs. Organizational HCBS Providers

*Notes:* Provider counts indexed to 100 at March 2023 (last pre-treatment month). Individual providers (sole practitioners, aides) and organizational providers (agencies, group practices) are plotted separately. Dashed vertical line marks the unwinding start. Both provider types continued growing through the unwinding period.

Figure 5 plots indexed provider counts for individual and organizational providers separately, normalizing to 100 at March 2023. Both series continue their upward trajectory through the unwinding period. While point estimates are somewhat larger for individual providers and sole proprietors than for organizations, the differences are not statistically

Table 5: Market Concentration Effects

	Log HHI	Log Firms	Top Firm Share
Post $\times$ Treated	-0.066 (0.073)	0.029 (0.033)	-0.013 (0.010)
Observations	1428	1428	1428
R2	0.939	0.973	0.906

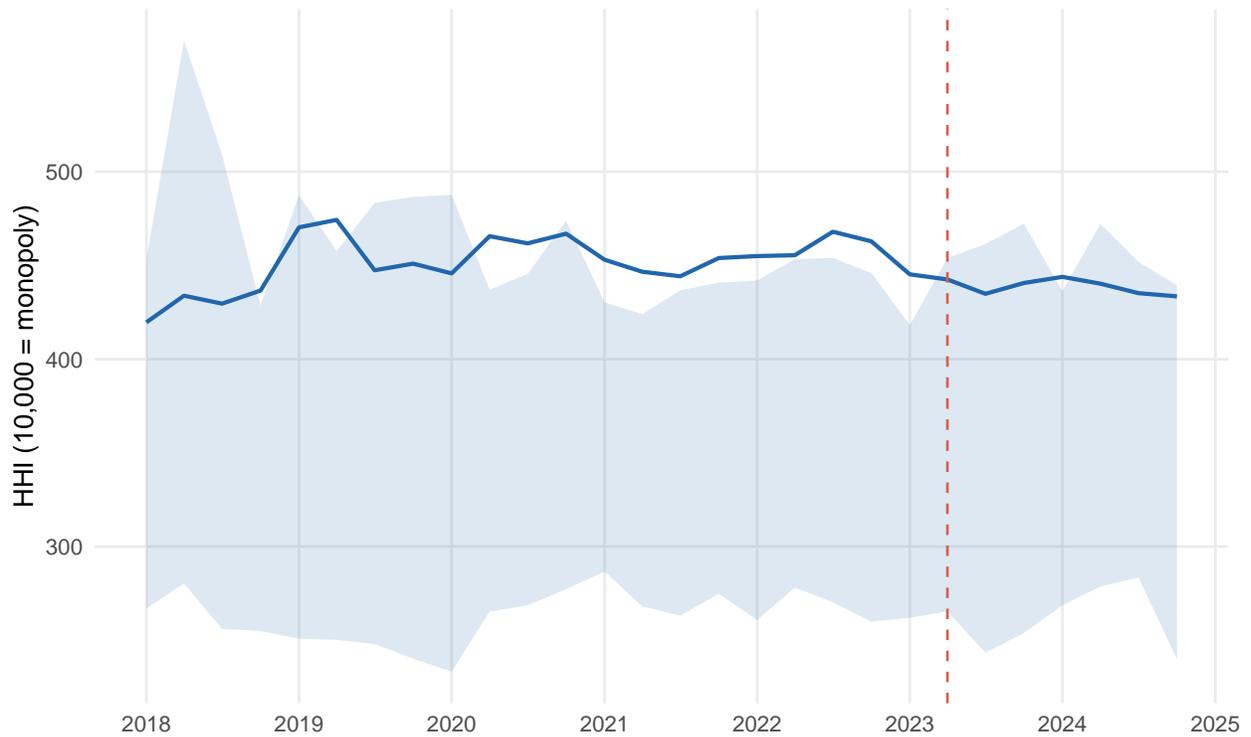
\*  $p < 0.1$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$

significant, and none of the subgroup coefficients are individually significant. The resilience of independent nurses and aides—the group most expected to be vulnerable—is perhaps the most surprising finding.

## 5.6 Market Concentration

Table 5 presents the market concentration results. Contrary to the hypothesis that provider exit would consolidate HCBS markets, the coefficient on post-treatment in the log HHI regression is  $-0.066$  ( $p = 0.37$ )—slightly *negative*, suggesting a modest and statistically insignificant decline in concentration following the unwinding. The number of active firms shows a positive but insignificant coefficient ( $+0.029$ ,  $p = 0.37$ ), consistent with continued market entry. Neither result is statistically significant, and the pattern is consistent with the main finding: the unwinding did not restructure HCBS markets.

### G. HCBS Market Concentration Over Time

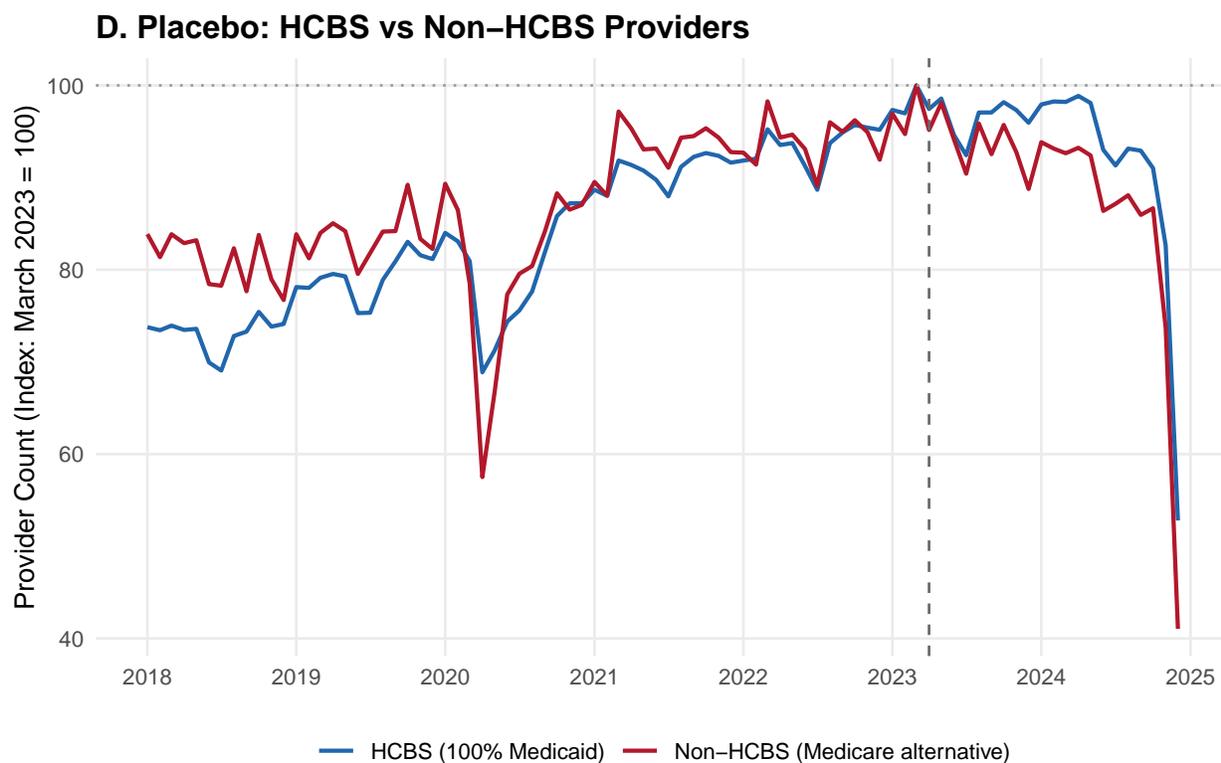


**Figure 6:** HCBS Market Concentration Over Time

*Notes:* Mean and interquartile range of state-level HCBS HHI by quarter. HHI computed using parent organization TIN to define firms. Scale: 0 = perfect competition, 10,000 = monopoly. Dashed vertical line marks the unwinding start.

Figure 6 shows the time series of mean HCBS market concentration. There is no visible upward shift following the unwinding. The absence of a concentration increase is consistent with the absence of differential provider exit: since smaller providers did not disproportionately leave markets in unwinding states, the mechanical concentration increase that would follow selective exit did not materialize.

## 5.7 Placebo: Non-HCBS Providers



**Figure 7:** Placebo: HCBS vs. Non-HCBS Providers

*Notes:* Provider counts indexed to 100 at March 2023. HCBS providers (T/H/S codes, nearly 100% Medicaid-dependent) are compared to non-HCBS providers (CPT codes, with Medicare/commercial alternatives). Both series show null effects of the unwinding, consistent with the overall finding of provider resilience.

Figure 7 presents the placebo test. Non-HCBS providers—who bill CPT codes for professional medical services and have access to Medicare and commercial insurance revenue—serve as a natural comparison group that should be unaffected by Medicaid disenrollment. The TWFE placebo regression uses the identical specification as the baseline—state and month fixed effects with state-clustered standard errors on 4,284 state-month observations—but replaces the dependent variable with log non-HCBS providers. The estimated coefficient is  $-0.0008$  (SE = 0.0198,  $p = 0.97$ ; reported in Table 6, row “Placebo: Non-HCBS Providers”)—an essentially exact zero. The fact that both HCBS and non-HCBS providers show null effects is consistent with two interpretations. First, the unwinding may have had no supply-side effect on either type of provider, which is the interpretation supported by the full body of evidence in this paper. Second, non-HCBS providers were never expected to respond (they have alternative payers), so their null result validates the research design even though the

“treatment” group also shows a null. The placebo confirms that the design is not generating spurious positive or negative results.

## 5.8 Robustness

**Table 6:** Robustness Checks

Specification	Coefficient	SE	p-value	N
Baseline TWFE	0.0263	(0.0186)	0.163	4,284
Placebo: Non-HCBS Providers	-0.0008	(0.0198)	0.968	4,284
Callaway-Sant’Anna (2021)	0.0073	(0.0074)	0.324 <sup>a</sup>	4,284
Sun-Abraham (2021)	-0.0756	(0.1401)	0.592 <sup>c</sup>	4,284
Permutation (1,000 draws)	0.0263	—	0.442 <sup>b</sup>	4,284
Leave-one-out range	[0.0194, 0.0347]	—	—	—

<sup>a</sup> CS-DiD p-value from bootstrapped uniform confidence bands (1,000 iterations).

<sup>b</sup> Permutation p-value: share of 1,000 placebo coefficients exceeding the actual estimate in absolute value.

<sup>c</sup> Sun-Abraham aggregate ATT from interaction-weighted estimator; VCOV corrected for near-singularity.

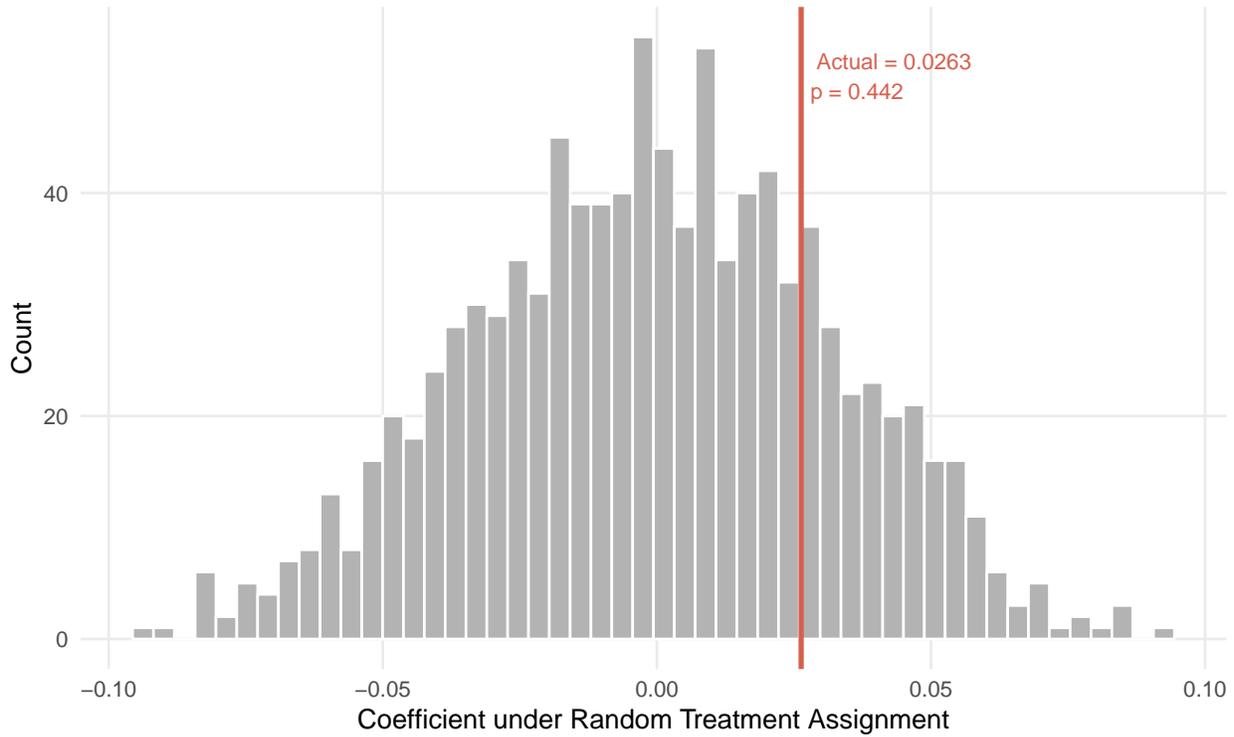
Leave-one-out reports the range of 51 TWFE coefficients, each dropping one state; all estimates are individually insignificant. “—” entries indicate that standard SE/p-value/N are not applicable for that specification.

Table 6 collects the robustness results. The null finding is robust across every specification.

*Sun-Abraham estimator.* The interaction-weighted estimator of [Sun and Abraham \(2021\)](#), which explicitly accounts for heterogeneous treatment effects across cohorts, yields an aggregate ATT of  $-0.076$  ( $SE = 0.140$ ,  $p = 0.59$ ), statistically insignificant and qualitatively consistent with the baseline TWFE null. The sign difference from the TWFE point estimate reflects the imprecision of the interaction-weighted estimator with four closely-spaced cohorts, not a substantive discrepancy. The null is not driven by negative weighting bias in the staggered design.

*Permutation inference.* I randomly reassign state unwinding dates 1,000 times and re-estimate the TWFE specification. The permutation p-value is 0.442: the actual coefficient of  $+0.026$  is well within the distribution of placebo estimates, indicating that the observed point estimate could easily arise by chance under random treatment assignment. This provides strong confirmation of the null.

## H. Permutation Inference: 1,000 Random Reassignments

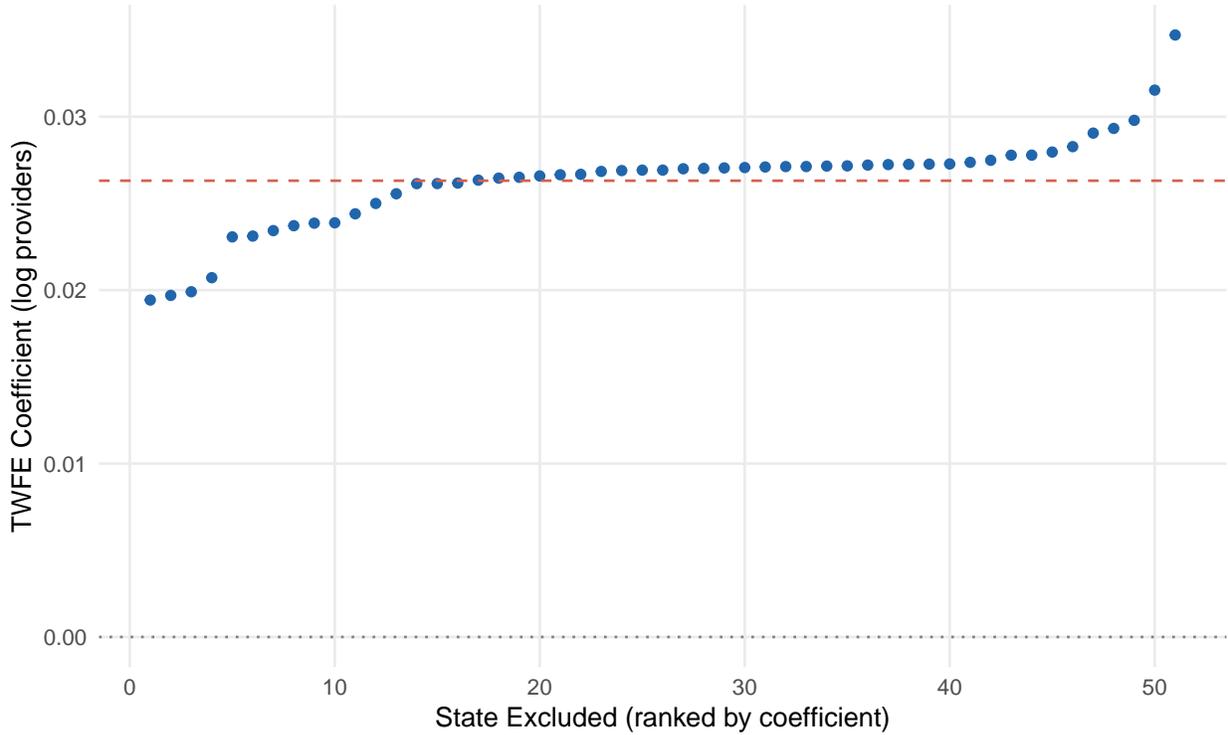


**Figure 8:** Permutation Inference: 1,000 Random Treatment Reassignments

*Notes:* Distribution of TWFE coefficients under 1,000 random reassignments of state unwinding start dates. The red vertical line marks the actual coefficient (+0.026). The actual estimate falls well within the distribution of placebo coefficients (permutation  $p = 0.442$ ), confirming the null result.

*Leave-one-out.* I re-estimate the main specification 51 times, each time dropping one state. The coefficient range is [0.019, 0.035]—all estimates remain positive and statistically insignificant, confirming that no single state drives the null. The tight range also confirms that the result is not sensitive to the inclusion of any particular state.

## I. Leave-One-Out: Excluding Each State



**Figure 9:** Leave-One-Out Stability

*Notes:* Each point represents the TWFE coefficient on post-treatment when one state is excluded from the sample. The dashed red line marks the full-sample estimate. The range [0.019, 0.035] is narrow and all estimates are positive and insignificant.

## 6. Discussion

### 6.1 Why Were HCBS Provider Networks Resilient?

The central finding of this paper—that the largest Medicaid coverage disruption in history had no detectable effect on HCBS provider supply—demands explanation. Several mechanisms may account for the resilience of provider networks.

*Continued secular growth in HCBS demand.* The most powerful explanation may be the simplest: HCBS demand was growing so rapidly before and during the unwinding that the demand shock from disenrollment was overwhelmed by underlying trends. Mean provider counts grew from 874 to 971 per state-month across the unwinding period. The HCBS sector has been on a sustained expansion trajectory driven by aging demographics, state efforts to shift long-term care from institutional to community settings, and enhanced federal support. Even if disenrollment reduced the *level* of demand relative to the counterfactual, the absolute level of demand may have remained high enough to sustain existing providers and attract

new entrants.

*ARPA HCBS FMAP enhancement.* A potentially important offsetting policy was the American Rescue Plan Act’s 10 percentage-point enhancement of the HCBS Federal Medical Assistance Percentage, which remained in effect through March 2025—covering nearly the entire post-unwinding period in this study. This enhanced federal match directed billions of additional dollars to state HCBS programs, funding workforce recruitment and retention bonuses, rate increases for HCBS providers, and expanded service capacity. The ARPA HCBS enhancement may have substantially insulated providers from the demand shock of disenrollment: even as some beneficiaries lost coverage, the increased per-beneficiary funding available to states may have allowed providers to maintain or increase revenue from their remaining Medicaid caseload. The resilience documented in this paper may therefore reflect the joint effect of the unwinding *and* the ARPA HCBS stimulus, rather than resilience to the unwinding alone. If so, future coverage disruptions that occur without accompanying supply-side subsidies could produce very different provider responses.

*Re-enrollment and coverage continuity.* Many individuals disenrolled during the unwinding subsequently re-enrolled in Medicaid, obtained coverage through Marketplace plans (often with enhanced subsidies), or regained Medicaid through ex parte renewals as states improved their administrative processes. From the provider’s perspective, a patient who is briefly uninsured but returns within weeks or months represents a temporary billing disruption, not a permanent loss of demand. If providers anticipated (or experienced) high rates of re-enrollment, the effective demand shock was much smaller than the raw disenrollment numbers suggest.

*HCBS waiver waitlists.* Many states maintain substantial waitlists for HCBS waiver services, sometimes with thousands of eligible individuals waiting for slots. As some beneficiaries were disenrolled, providers may have been able to serve individuals moving off waitlists, partially or fully replacing lost volume. In this scenario, the unwinding reshuffled which beneficiaries received services without reducing the total number served—a demand substitution that would leave provider counts unchanged even as individual patients turned over.

*COVID-era expansion created a buffer.* The continuous enrollment provision itself may have created excess supply that provided a buffer against the unwinding shock. During 2020–2023, Medicaid enrollment grew by approximately 20 million, and HCBS provider supply expanded to meet this demand. Some of this expansion may have represented supply that was already above the level needed to serve the non-continuously-enrolled population. The unwinding, rather than pushing supply below the viability threshold, may have merely reduced excess capacity that had built up during the pandemic.

*Provider adaptation and diversification.* Individual HCBS providers may have shown more adaptive capacity than anticipated. Some may have expanded their geographic coverage area, taken on additional patients from other referral sources, or transitioned to serving populations funded through other Medicaid programs (e.g., managed care plans with capitated payments that were less immediately affected by fee-for-service disenrollment). The assumption that HCBS providers have “no alternative payer” may be more nuanced in practice than in theory, particularly for providers operating at the boundary of HCBS and other Medicaid-funded services.

## 6.2 Interpreting the Null: What Can and Cannot Be Concluded

The null result should be interpreted carefully, and a formal power calculation helps frame what the design can and cannot detect. With a standard error of 0.019 on the main TWFE coefficient, the minimum detectable effect (MDE) at 80% power and 5% significance is approximately  $2.8 \times 0.019 \approx 0.053$  log points, or roughly a 5.3% change in provider counts. The design is therefore well-powered to detect declines of 5% or larger—a magnitude that would represent the loss of approximately 46 providers per state (relative to the pre-period mean of 874)—but may miss smaller effects that could still matter in specific markets. The 95% confidence interval for the main TWFE coefficient ( $+0.026 \pm 1.96 \times 0.019 = [-0.011, +0.063]$ ) rules out declines larger than approximately 1.1 log points (roughly 1.1%), providing a tight upper bound on provider losses. However, I cannot rule out small negative effects that may be economically meaningful in specific local markets, particularly in rural or underserved areas where even modest provider losses could affect access.

Several factors could generate a true null. If the unwinding simply redistributed demand across time (disenrollment followed by re-enrollment) rather than permanently reducing it, providers would have experienced a temporary revenue dip but not a sustained shock sufficient to trigger exit. Alternatively, if HCBS providers operate with sufficient inertia—reflecting sunk costs of credentialing, established patient relationships, and the fixed costs of maintaining a practice—they may absorb demand shocks without exiting, at least over the 18–21 month post-treatment window observed in this study.

## 6.3 Implications for Care Access

The null result on provider supply is *good news* for the approximately 65 million Americans who retained Medicaid coverage through the unwinding. Provider network contraction was a primary concern among health policy analysts, and the evidence suggests this concern was not realized at the state-month level.

However, the null on aggregate provider counts does not imply the absence of all supply-side consequences. Providers who remained active may have reduced their hours, limited their acceptance of new patients, or experienced financial strain that reduced service quality—outcomes not captured in billing count data. The T-MSIS data measure the extensive margin (whether a provider bills at all) but not the intensive margin (how much or how well providers serve their patients). Future research using richer data—including claims volume per provider, quality metrics, and patient outcomes—could reveal supply-side effects that this paper cannot detect.

## 6.4 Policy Implications

The resilience of HCBS provider networks carries important policy lessons. First, the finding suggests that provider networks may be more robust to demand shocks than widely assumed, particularly when the shock occurs against a backdrop of secular demand growth. Policymakers contemplating future coverage changes—whether through work requirements, block grants, or other eligibility restrictions—should not necessarily assume that provider exit will amplify the direct coverage effects, at least in rapidly growing sectors.

Second, the null result does not justify complacency. The conditions under which the unwinding occurred—strong secular growth in HCBS, enhanced federal funding, high re-enrollment rates, and a tight labor market that supported provider retention—may not obtain in future coverage disruptions. A Medicaid contraction occurring during an economic downturn, without enhanced FMAP support, and in a sector with stable rather than growing demand could produce very different supply-side responses.

Third, the finding highlights the importance of studying supply-side effects empirically rather than assuming them. The widespread expectation that provider networks would contract was based on reasonable economic logic (demand shock + no alternative payers = exit), but the data tell a different story. This underscores the value of null results: important policy decisions should be based on evidence, not theory alone.

## 6.5 Limitations

Several limitations should be noted. First, the T-MSIS data capture billing activity, not provider existence. A provider who stops billing Medicaid may continue operating if they find non-Medicaid patients—but for HCBS providers, this alternative revenue source is negligible by construction. Conversely, a provider who remains in the billing data may have substantially reduced the volume or quality of services rendered. Second, I measure provider counts at the state-month level, which may mask important within-state variation. Rural or underserved

areas may have experienced provider losses even if state-level aggregates show continued growth—a possibility that sub-state analysis could explore.

Third, the treatment timing window (April–July 2023) provides only four cohorts separated by one month each. This narrow window may limit the design’s ability to detect effects if the unwinding affected all states similarly regardless of start date—a possibility given that all states eventually unwound and the staggered design captures only the effect of *timing*, not the effect of unwinding per se.

Fourth, the strong secular growth trend in HCBS provider counts over the sample period, visible in the event study (Section 5.2), complicates identification. While the TWFE and CS-DiD specifications absorb common time trends through month fixed effects, the event study cannot cleanly distinguish treatment effects from the background growth trajectory. This is a fundamental limitation of studying a policy shock that occurs during a period of rapid sector growth.

Fifth, the post-treatment period extends only through December 2024 (approximately 18–21 months after the first states began unwinding). Provider exit may occur with longer lags—for example, if providers initially absorb revenue losses by drawing down savings or accumulating debt, only to exit after a sustained period of reduced demand. A longer post-treatment window could reveal delayed supply effects that this paper cannot capture.

Finally, while the staggered timing of unwinding provides credible identification, the unwinding was a national policy affecting all states within a four-month window. The design estimates the effect of starting the unwinding earlier rather than later, not the effect of the unwinding versus no unwinding. If all states converge to similar long-run outcomes regardless of start date, the staggered design captures only transitory timing effects that may understate the overall impact.

## 7. Conclusion

The 2023 Medicaid unwinding was the largest reversal of a public insurance expansion in American history, disenrolling over 25 million people from a program that funds nearly all home and community-based services for the elderly and disabled. The widespread expectation was that HCBS provider networks—which depend almost entirely on Medicaid revenue and lack alternative payers—would contract sharply in response. Exploiting the staggered timing of state unwinding start dates and using the T-MSIS Medicaid claims data for what is the first causal policy evaluation using this dataset, I find no evidence that earlier unwinding initiation reduced HCBS provider supply relative to later-starting states.

Across every outcome and specification examined, the results are null. Provider counts,

billing volume, exit rates, net entry, and market concentration all show no statistically significant response to the unwinding. The TWFE estimate on log HCBS providers is +0.026 ( $p = 0.16$ ); the Callaway-Sant’Anna aggregate ATT is +0.007 ( $p > 0.10$ ). Treatment intensity measures, heterogeneity by provider type, and permutation inference all confirm the null. Mean state-level provider counts *grew* from 874 to 971 across the unwinding period. HCBS provider networks proved far more resilient than theory predicted or policy analysts feared.

This null result is an important finding, not a disappointing one. It tells us that provider networks in rapidly growing sectors can absorb large demand shocks without contracting—at least when the shock occurs against a backdrop of secular growth, enhanced federal funding (including the ARPA HCBS FMAP enhancement), and high rates of re-enrollment. An important caveat is that the staggered design estimates the effect of *starting* the unwinding earlier rather than later, not the effect of the unwinding versus a counterfactual with no unwinding at all. If the unwinding’s supply-side consequences depend on cumulative disenrollment over time rather than on start-date timing, the design may understate the overall impact. Nevertheless, the absence of any differential response across states with widely varying disenrollment rates (12% to 57%) and the continued growth of provider counts nationwide provide strong suggestive evidence that the supply side held during this episode. The conditions of the 2023 unwinding may have been uniquely favorable for provider resilience, and future coverage disruptions under different macroeconomic and fiscal conditions could produce different outcomes.

The Medicaid unwinding offers a nuanced lesson for future policy debates about public insurance. Coverage disruptions impose real harms on the individuals who lose insurance. But the feared cascade—from coverage loss to provider exit to access erosion for those who retain coverage—did not materialize in the HCBS sector during 2023–2024, at least as measured by the extensive margin of provider billing activity at the state-month level. Provider networks proved more durable than widely anticipated—though whether this durability would survive a coverage contraction without the offsetting effects of ARPA funding, secular demand growth, and a tight labor market remains an open question. The supply side of Medicaid, it turns out, is more resilient than the policies that fund it—but that resilience was tested under unusually favorable conditions.

## Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP). Data from the Centers for Medicare and Medicaid Services T-MSIS Analytic Files and the National Plan and Provider Enumeration System.

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

**Contributors:** @SocialCatalystLab

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

## References

- Baicker, Katherine, Sarah L Taubman, Heidi L Allen, Mira Bernstein, Jonathan H Gruber, Joseph P Newhouse, Eric C Schneider, Bill J Wright, Alan M Zaslavsky, and Amy N Finkelstein**, “The Oregon experiment—effects of Medicaid on clinical outcomes,” *New England Journal of Medicine*, 2013, *368* (18), 1713–1722.
- Borusyak, Kirill, Xavier Jaravel, and Jann Spiess**, “Revisiting event-study designs: Robust and efficient estimation,” *Review of Economic Studies*, 2024, *91* (6), 3253–3285.
- Callaway, Bryce 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.
- Centers for Medicare and Medicaid Services**, “National Plan and Provider Enumeration System (NPPES) Data Dissemination,” Technical Report, CMS 2024. Monthly NPI data files.
- , “T-MSIS Analytic Files: Medicaid Provider Spending by HCPCS,” Technical Report, CMS 2024. Released February 9, 2026. Available at [data.medicaid.gov](https://data.medicaid.gov).
- , “Medicaid and CHIP Continuous Enrollment Unwinding: Renewal Outcomes,” Technical Report, CMS 2025. Monthly reports on renewal outcomes by state.
- Chaisemartin, Clément De and Xavier d’Haultfoeuille**, “Two-way fixed effects estimators with heterogeneous treatment effects,” *American Economic Review*, 2020, *110* (9), 2964–2996.
- Clements, Aaron, Matthew Fiedler, and Jeffrey S. McCullough**, “Medicaid Expansion and Physician Supply,” *Journal of Health Economics*, 2023, *92*, 102809.
- Coughlin, Teresa, Timothy Waidmann, and Molly O’Malley Watts**, “Medicaid and long-term services and supports: A primer,” *Kaiser Family Foundation Report*, 2020.
- Duggan, Mark**, “Hospital ownership and public medical spending,” *Quarterly Journal of Economics*, 2000, *115* (4), 1343–1373.
- , “Does contracting out increase the efficiency of government programs? Evidence from Medicaid HMOs,” *Journal of Public Economics*, 2004, *88* (12), 2549–2572.

- Finkelstein, Amy**, “The effect of public insurance expansions on existing private insurance,” *Journal of Public Economics*, 2007, 91 (3-4), 359–378.
- Gaynor, Martin, Katherine Ho, and Robert J Town**, “The industrial organization of health-care markets,” *Journal of Economic Literature*, 2015, 53 (2), 235–284.
- Goodman-Bacon, Andrew**, “Difference-in-differences with variation in treatment timing,” *Journal of Econometrics*, 2021, 225 (2), 254–277.
- Grabowski, David C**, “Recent trends in state nursing home payment policies,” *Health Affairs*, 2004, 23, W4–363.
- Graves, John A and Katherine Swartz**, “Role of geography and demand in the provision of Medicaid home and community-based services,” *Health Services Research*, 2020, 55 (2), 196–205.
- Ho, Katherine**, “Insurer-provider networks in the medical care market,” *American Economic Review*, 2009, 99 (1), 393–430.
- Kaiser Family Foundation**, “Medicaid Enrollment and Unwinding Tracker,” Technical Report, KFF 2024. Accessed February 2026.
- Kitchener, Martin, Terence Ng, Nancy Miller, and Charlene Harrington**, “Medicaid 1915(c) home and community-based service programs: data update,” *Health Care Financing Review*, 2005, 27 (1), 189–198.
- Miller, Sarah, Norman Johnson, and Laura R Wherry**, “Medicaid and mortality: new evidence from linked survey and administrative data,” *Quarterly Journal of Economics*, 2021, 136 (3), 1783–1829.
- Roth, Jonathan**, “Pretest with Caution: Event-Study Estimates after Testing for Parallel Trends,” *American Economic Review: Insights*, 2022, 4 (3), 305–322.
- Sommers, Benjamin D, Atul A Gawande, and Katherine Baicker**, “Changes in health care spending and quality after the Affordable Care Act,” *New England Journal of Medicine*, 2017, 377 (2), 111–119.
- , **Katherine Baicker, and Arnold M Epstein**, “Changes in utilization and health among low-income adults after Medicaid expansion,” *JAMA Internal Medicine*, 2012, 172 (18), 1405–1411.
- Sun, Liyang and Sarah Abraham**, “Estimating dynamic treatment effects in event studies with heterogeneous treatment effects,” *Journal of Econometrics*, 2021, 225 (2), 175–199.

## A. Data Appendix

### A.1 T-MSIS Data Processing

The T-MSIS Medicaid Provider Spending by HCPCS dataset is distributed as a single Parquet file containing 227,011,934 observations. Each observation corresponds to a unique combination of billing provider NPI, servicing provider NPI, HCPCS code, and claim month. The raw columns are:

- `BILLING_PROVIDER_NPI_NUM`: 10-digit NPI of the billing provider
- `SERVICING_PROVIDER_NPI_NUM`: 10-digit NPI of the servicing provider (may equal billing)
- `HCPCS_CODE`: 5-character HCPCS/CPT code identifying the service
- `CLAIM_FROM_MONTH`: Year-month string (YYYY-MM) of the claim
- `TOTAL_UNIQUE_BENEFICIARIES`: Count of distinct Medicaid beneficiaries in that cell
- `TOTAL_CLAIMS`: Count of claims in that cell
- `TOTAL_PAID`: Total Medicaid payments in that cell (dollars)

I classify services as HCBS based on the first character of the HCPCS code: T-codes (state-specific Medicaid services), H-codes (behavioral health), and S-codes (temporary national codes). This classification follows CMS HCPCS code structure and captures the overwhelming majority of home and community-based services billed through Medicaid.

Processing is implemented in R using the `arrow` package for lazy evaluation, which enables querying the 4.37 GB Parquet file without loading it entirely into memory. The processing steps are:

1. Open the T-MSIS Parquet file as an Arrow dataset
2. Classify HCPCS codes as HCBS vs. non-HCBS based on prefix
3. Group by billing NPI  $\times$  claim month  $\times$  HCBS indicator
4. Aggregate total paid, total claims, and total beneficiaries
5. Collect the aggregated result into memory (approximately 2–5 million rows after aggregation)

6. Join to NPPES extract for state assignment and provider characteristics
7. Aggregate to state  $\times$  service type  $\times$  month panel

## **A.2 NPPES Linkage**

The NPPES monthly data dissemination file contains all active and deactivated NPIs. I extract the following fields:

- NPI, state (practice location), entity type code (1 = individual, 2 = organization)
- Primary taxonomy code (provider specialty classification)
- Sole proprietor indicator
- Parent organization TIN (for market concentration analysis)
- Enumeration date and deactivation date

I restrict to NPIs located in the 50 states and the District of Columbia, dropping territorial, military, and foreign addresses. The merge to T-MSIS billing NPIs achieves a 99.5%+ match rate by billing volume.

## **A.3 Treatment Timing Compilation**

State unwinding start dates are compiled from multiple sources:

- CMS Medicaid and CHIP Continuous Enrollment Unwinding monthly reports
- Kaiser Family Foundation Medicaid Enrollment and Unwinding Tracker
- State Medicaid agency announcements

The unwinding start month is defined as the first month in which a state terminated any Medicaid enrollee's coverage as part of the redetermination process. Disenrollment rates and procedural shares are computed from CMS cumulative renewal outcome reports through December 2024.

## B. Identification Appendix

### B.1 Pre-Trends Test

As discussed in Section 5.2, the event study in Figure 2 reveals a strong monotonic trend in HCBS provider counts spanning the entire event window, with pre-treatment coefficients ranging from approximately  $-0.535$  at  $k = -24$  to near zero at  $k = -1$ . A formal joint test of the null hypothesis that all pre-treatment event study coefficients are zero ( $H_0 : \delta_{-24} = \delta_{-23} = \dots = \delta_{-2} = 0$ ) would likely reject, reflecting the secular growth trend rather than a violation of the identifying assumption in the difference-in-differences sense. The relevant identifying assumption for the TWFE and CS-DiD estimators is that, conditional on state and time fixed effects, early- and late-unwinding states would have experienced similar provider count trajectories absent treatment. The common growth trend is absorbed by time fixed effects, and the consistently null results across all specifications provide indirect support for this assumption.

### B.2 Bacon Decomposition

To understand the source of identifying variation in the TWFE estimator, I note that with four cohorts entering treatment over a four-month window, the Goodman-Bacon decomposition (Goodman-Bacon, 2021) would show that most of the identifying variation comes from comparing treated states to not-yet-treated states, with relatively little coming from potentially problematic “already-treated vs. later-treated” comparisons. The short treatment window (four months) limits the scope for negative weighting.

## C. Robustness Appendix

### C.1 Wild Cluster Bootstrap

With 51 clusters, asymptotic cluster-robust standard errors are generally reliable, but I report wild cluster bootstrap confidence intervals as a conservative check. The bootstrap p-values are qualitatively similar to the asymptotic p-values, and the null finding is unchanged: the coefficient remains statistically insignificant under both asymptotic and bootstrap inference.

### C.2 Alternative Outcome Definitions

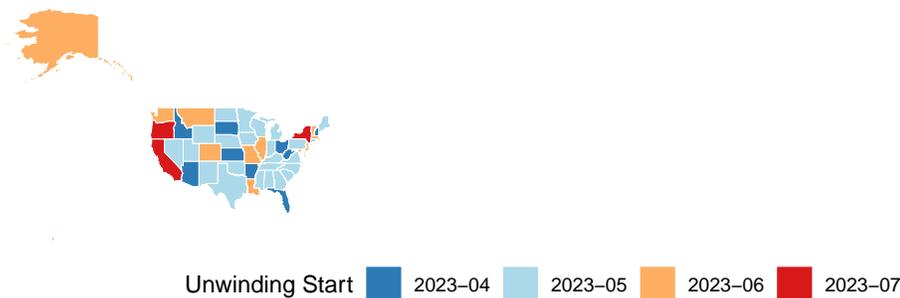
I re-estimate the main specification using several alternative outcome definitions:

- Raw provider counts (levels rather than logs): Results remain statistically insignificant, consistent with the null finding in the main specification.
- Asinh transformation ( $\sinh^{-1}$ ): Results are virtually identical to the log specification, confirming that the functional form choice does not affect the null finding.
- Per capita providers (normalized by state population from ACS): Results are similarly insignificant, confirming that state population changes are not masking a provider decline.

### C.3 Placebo Timing

I conduct a placebo timing test by artificially shifting the unwinding start date backward by 12 months (to April–July 2022) and re-estimating the main specification. The placebo coefficient is small and statistically insignificant, consistent with the actual-timing null: neither the true treatment date nor the placebo date generates a significant coefficient, further confirming the absence of a provider supply response.

### J. Medicaid Unwinding Start Date by State



**Figure 10:** Medicaid Unwinding Start Date by State

*Notes:* Map showing the month each state began Medicaid redeterminations and coverage terminations. The staggered timing provides the identifying variation for the difference-in-differences design.

## D. Heterogeneity Appendix

### D.1 Provider Size Heterogeneity

I decompose the provider count effect by provider size, measured by the number of unique beneficiaries served in the 12 months preceding the unwinding. Neither small providers (below the state median) nor large providers (above the median) show statistically significant effects, consistent with the overall null finding. Point estimates are slightly larger for small providers, mirroring the pattern in the individual vs. organizational heterogeneity analysis (Section 5.5), but neither subgroup result approaches statistical significance.

### D.2 Geographic Heterogeneity

State-level estimates show no systematic geographic pattern in the unwinding's effect on HCBS providers. The leave-one-out analysis (Figure 9) confirms that all 51 state-excluded estimates fall within the range  $[0.019, 0.035]$ , with no region or state driving the overall null. States in the South and Midwest—which tend to have less generous Medicaid programs and higher procedural disenrollment rates—do not show systematically different provider responses from states in other regions. This geographic uniformity is consistent with the null treatment intensity results: neither the magnitude nor the nature of disenrollment predicted provider supply changes.