

@ai1scl

Paying More, Getting Less? The Perverse Effects of Medicaid HCBS Reimbursement Rate Increases on Provider Supply

APEP Autonomous Research* APEP Working Paper 0341 v1, incorporating feedback from

February 18, 2026

Abstract

Medicaid home and community-based services face chronic provider shortages. States raise reimbursement rates to attract providers, yet evidence on effectiveness is thin. Using provider-level claims from T-MSIS covering 52 jurisdictions over 84 months (2018–2024), I identify 20 jurisdictions with sustained personal care rate increases of 15% or more and estimate their causal effect on provider supply. Both TWFE and Callaway-Sant’Anna estimators find no evidence that rate increases expand the HCBS workforce. Point estimates are negative and statistically insignificant under asymptotic and wild cluster bootstrap inference. Separating ARPA-era cohorts (post April 2021) from pre-ARPA cohorts yields consistent nulls. Mechanism tests find no significant change in organizational billing share, though individual providers show larger declines than organizations. These findings suggest Medicaid’s HCBS provider supply crisis cannot be solved by rate increases alone.

JEL Codes: I11, I13, I18, J44, H75

Keywords: Medicaid, HCBS, reimbursement rates, provider supply, home care, ARPA, difference-in-differences

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

1. Introduction

In 2024, over 800,000 Americans sat on waiting lists for Medicaid home and community-based services (?). These services—personal care assistance, attendant care, home health aides—allow elderly and disabled individuals to live in their communities rather than institutions. Yet the workforce that delivers them is in crisis. Provider turnover exceeds 50% annually, vacancy rates top 20% in many states, and the pipeline of new entrants is shrinking (?). The conventional policy response is straightforward: raise reimbursement rates. If Medicaid pays more, more providers should enter the market and existing providers should expand capacity. Between 2021 and 2023, the American Rescue Plan Act channeled \$37 billion into home care, with rate increases the most common investment (?). But does paying more actually get more?

This paper provides causal evidence on the effect of Medicaid reimbursement rate increases on HCBS provider participation, capacity, and service volume using provider-level claims data. I exploit staggered rate increases across 20 jurisdictions during 2018–2024, detected directly from payment data in the Transformed Medicaid Statistical Information System (T-MSIS). The T-MSIS dataset, published by HHS in February 2026, contains the near-universe of Medicaid provider billing records—227 million claim-month observations covering all 50 states plus territories—disaggregated by billing provider, procedure code, and month. Linked to the National Plan and Provider Enumeration System (NPPES) for geographic identification, this data enables construction of a comprehensive state-level monthly panel of HCBS provider supply, billing volume, and payment rates.

My identification strategy exploits variation in the timing and magnitude of state-level personal care reimbursement rate increases. I measure rate changes from the T-MSIS data by detecting sustained jumps of 15% or more in average payment per claim for personal care procedure codes (T1019, T1020, S5125, S5130). Twenty states exhibit such jumps, with most increases between 17% and 260% and one extreme outlier (Wyoming at 1,422%), while 32 jurisdictions serve as never-treated controls. Many of the post-2021 increases were funded by the American Rescue Plan Act’s 10% FMAP enhancement for HCBS, providing an institutional narrative for the exogeneity of treatment timing: federal legislation created a funding shock, and states implemented rate increases at different times based on legislative calendars, CMS approval processes, and administrative capacity rather than contemporaneous provider supply conditions.

The main finding is that there is no evidence rate increases expand the personal care provider workforce. The TWFE estimate for log provider counts is negative and statistically insignificant under both clustered standard errors and wild cluster bootstrap inference (??)—

with the 95% confidence interval spanning $[-0.50, 0.37]$, we can rule out large positive supply responses but not modest effects in either direction. The Callaway-Sant’Anna heterogeneity-robust ATT (?) confirms the null with attenuated magnitudes. Both total claims and beneficiaries served show similar patterns. Total Medicaid spending increases mechanically—existing claims are paid at higher rates—but the volume of services delivered does not respond.

The null is robust across a comprehensive battery of tests. Placebo analyses confirm that personal care rate increases have no effect on unrelated services: E/M office visit providers show null effects. A pre-treatment lead test (placebo treatment shifted 12 months early) confirms the absence of anticipation effects. Alternative detection thresholds (10%, 20%, 25%), median-based rate detection, and varying rolling-mean bandwidths (2, 3, 4, 6 months) yield consistent results. Excluding the COVID-19 onset period does not change the estimates. State-specific linear trends do not alter the conclusion. Crucially, I separate ARPA-era treatment cohorts (post April 2021, 7 states) from pre-ARPA cohorts—the ARPA subsample, which benefits from stronger exogeneity due to the federal funding shock, yields similarly null results.

Heterogeneity analyses reveal a striking pattern. Individual providers—sole practitioners who constitute the majority of the direct care workforce—show a suggestive *decline* following rate increases, though imprecisely estimated. Organizational providers decline by a similar magnitude but with more noise. However, mechanism tests do not confirm the consolidation hypothesis: the share of Type 2 (organizational) billing NPIs shows no statistically significant change in treated states relative to controls after rate hikes ($\hat{\beta} = -0.0018, p = 0.84$), and average claims per provider declines insignificantly ($\hat{\beta} = -0.20, p = 0.15$). The null supply response is thus pervasive across both the extensive and intensive margins.

These findings contribute to three literatures. First, I advance the economics of Medicaid provider supply by providing quasi-experimental estimates of the reimbursement rate elasticity of HCBS provider participation. Prior work has documented cross-sectional correlations between Medicaid fee-for-service rates and physician participation (???), and evidence suggests that home care wages affect worker entry and turnover (???). But the HCBS workforce—predominantly paraprofessional, minimally credentialed, and heavily female—operates under fundamentally different labor market conditions than physicians. Low wages coexist with employer-side market concentration (?), and occupational licensing barriers interact with wage incentives in complex ways (?). My estimates suggest the short-run supply elasticity for personal care providers is close to zero, consistent with structural barriers to entry dominating price incentives.

Second, I contribute to the evaluation of ARPA pandemic-era spending. The \$37 billion

in HCBS investment was the largest single federal investment in home care in American history (?). My results suggest that the rate-increase component of this investment—the most popular state strategy—did not achieve its stated goal of expanding provider networks. This does not mean the money was wasted; providers who remained received higher compensation, which may reduce turnover in the long run. But the headline supply response was zero. The ARPA-era subsample analysis isolates treatment cohorts with the strongest claim to federal exogeneity and still finds no positive effect.

Third, I demonstrate the research value of the T-MSIS provider spending data for causal inference. Prior Medicaid research has relied on aggregate state expenditure reports (CMS-64), enrollment files, or survey data with small samples and long publication lags (?). The T-MSIS data enables real-time measurement of provider behavior at granular levels—individual NPIs billing specific procedure codes each month—opening a new frontier for Medicaid program evaluation. This paper is among the first to exploit this data for causal analysis.

2. Institutional Background

2.1 Medicaid Home and Community-Based Services

Medicaid is the primary payer for long-term services and supports (LTSS) in the United States, financing approximately 60% of all nursing home care and a growing share of home and community-based alternatives (?). HCBS encompasses a wide range of services—personal care assistance, attendant care, habilitation, respite care, home health, and behavioral health—that enable individuals with disabilities and the elderly to live in community settings rather than institutions.

The shift from institutional to community-based care has been a central goal of Medicaid policy since the *Olmstead v. L.C.* Supreme Court decision in 1999 mandated community integration for individuals with disabilities. By 2020, HCBS accounted for 57% of Medicaid LTSS spending, up from 30% in 2000 (?). Yet this rebalancing has been constrained by persistent workforce shortages. The HCBS workforce—direct care workers, personal care aides, home health aides—is one of the largest and lowest-paid occupational categories in the American economy, with median hourly wages below \$15 and annual turnover rates exceeding 50% (?).

The HCBS provider market differs fundamentally from physician services in several important respects. Providers are disproportionately female (87%), non-white (62%), and lacking post-secondary credentials (?). Many are sole proprietors or independent contractors. Entry barriers are low but non-trivial: most states require training (typically 40–120 hours), background checks, and enrollment in the Medicaid provider system. Exit barriers are equally

low—providers can simply stop billing. The resulting market exhibits extreme dynamism: analysis of the T-MSIS data reveals that only 6% of HCBS providers bill continuously across the 2018–2024 period, and the median provider tenure is just 22 months. This churning is consistent with prior evidence on direct care worker labor mobility (?) and the retention challenges documented in San Francisco’s living wage experiment (?).

2.2 Reimbursement Rate Setting

States set Medicaid reimbursement rates for HCBS through fee schedules that specify payment amounts for each procedure code. Personal care services are typically billed in 15-minute increments using HCPCS codes T1019 (personal care services) and S5125 (attendant care), with rates varying dramatically across states. In the pre-ARPA period, average payment per claim for personal care codes ranged from under \$3 per unit in some states to over \$500 in others, reflecting differences in service intensity definitions, billing practices, and cost-of-living adjustments.

Rate-setting is administrative and discretionary. Unlike Medicare, which uses nationally standardized fee schedules updated by formula, Medicaid rates are set by each state independently, subject only to the federal requirement that rates be “sufficient to enlist enough providers so that care and services are available” (Section 1902(a)(30)(A) of the Social Security Act). In practice, many states have not updated HCBS rates for years or decades, allowing real rates to erode through inflation. The adequacy of Medicaid payment rates for attracting and retaining providers has been a long-standing concern across service categories (???)

Rate increases typically require state legislative action or administrative rulemaking, creating natural delays between the decision to raise rates and implementation. This administrative lag generates variation in treatment timing that is plausibly exogenous to contemporaneous provider supply conditions—the timing reflects bureaucratic and legislative processes rather than real-time labor market adjustment.

2.3 The ARPA HCBS Funding Shock

The American Rescue Plan Act of 2021 (ARPA), enacted March 11, 2021, included Section 9817, which provided states with a temporary 10 percentage point increase in the Federal Medical Assistance Percentage (FMAP) for HCBS expenditures incurred between April 1, 2021, and March 31, 2022 (?). The enhanced FMAP effectively gave states additional federal dollars for every HCBS dollar spent, generating approximately \$37 billion in savings that states were required to reinvest in HCBS improvements.

States submitted spending plans to CMS detailing how they would use the ARPA funds. The most common investment category was provider payment increases: over 40 states included rate increases in their spending plans (?). However, implementation varied dramatically across states in terms of timing, magnitude, permanence, and breadth. Some states implemented rate increases within months of ARPA’s passage (e.g., Indiana in 2021 Q2), while others took more than two years (e.g., Ohio in 2023 Q4). Rate increases ranged from modest (15–20%) to dramatic: Mississippi raised personal care rates by 162%, Ohio by 228%, and Arizona by 254%. Some rate increases were permanent fee schedule changes; others were temporary bonuses that expired when ARPA funds were exhausted. Some states raised rates for all HCBS services; others targeted specific service categories.

This heterogeneity in timing, magnitude, and scope provides the identifying variation for my analysis. Crucially, the federal legislation itself was exogenous to individual state provider supply conditions—all states received the same FMAP enhancement regardless of their workforce situation. The variation in *when* and *how much* states raised rates was driven by state-specific legislative calendars, CMS approval processes, and administrative capacity.

3. Conceptual Framework

A simple model of provider participation in Medicaid HCBS clarifies the expected effects of rate increases and the conditions under which they might fail.

Consider a potential provider i who chooses between entering the Medicaid HCBS market and an outside option (other employment, non-Medicaid home care, or not working). Let r_s denote the Medicaid reimbursement rate in state s , c_i the provider’s cost of service delivery (including training, background checks, transportation, and opportunity cost), and w_i the outside option wage. Provider i enters if:

$$r_s \cdot q_i - c_i > w_i \tag{1}$$

where q_i is the expected number of billable units per period. The supply of providers in state s is:

$$N_s = \int \mathbb{I}[r_s \cdot q_i - c_i > w_i] dF(c_i, w_i, q_i) \tag{2}$$

A rate increase $\Delta r > 0$ shifts the participation threshold, and the supply response depends on the density of marginal providers near the threshold:

$$\frac{\partial N_s}{\partial r_s} = \int q_i \cdot f(r_s \cdot q_i - w_i) dF(w_i, q_i) \geq 0 \tag{3}$$

This derivative is unambiguously non-negative in the standard model: higher rates should weakly increase provider supply. However, several mechanisms could generate a zero or negative empirical relationship:

Prediction 1: Near-zero supply elasticity. If the distribution of $c_i - w_i$ has thin tails near the participation threshold, few marginal providers exist to be attracted by a rate increase. This can occur when entry barriers (training, licensing, background checks) create a discrete cost of entry that dominates the rate signal. The supply curve is locally inelastic.

Prediction 2: Offsetting general equilibrium effects. Rate increases for Medicaid may simultaneously increase the outside option if they raise wages in the broader home care market. If competing employers (private-pay home care, assisted living facilities) match Medicaid rate increases, the net effect on relative compensation is zero. This is consistent with evidence on monopsony dynamics in care labor markets (?) and quit behavior in response to wage changes (?).

Prediction 3: Organizational consolidation. If rate increases primarily benefit organizational providers (who bill on behalf of employed workers), they may accelerate consolidation—organizations expand while individual providers exit. The net effect on NPI counts could be zero or negative even as total capacity (hours delivered) increases.

Prediction 4: Reverse causality. If states raise rates precisely when provider supply is declining (responsive policy), the estimated effect conflates the rate increase with the pre-existing decline. This is an identification concern rather than a mechanism, addressed through event-study designs.

My empirical analysis tests these predictions by examining (1) whether the aggregate supply elasticity is significantly positive, (2) whether heterogeneity across provider types is consistent with consolidation, (3) whether mechanism tests directly confirm organizational share shifts, and (4) whether event-study patterns are consistent with reverse causality or clean identification.

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, published by the Department of Health and Human Services in February 2026. This dataset contains the near-universe of Medicaid provider billing records at the billing provider \times servicing provider \times HCPCS code \times month level. For each observation, the data report total unique beneficiaries, total claims, and total amount paid. The data cover all 50 states, the District of Columbia, and U.S. territories from January 2018 through

December 2024—84 months of coverage encompassing 227 million observations and \$1.09 trillion in cumulative payments. Of these, 52 jurisdictions (the 50 states, DC, and the U.S. Virgin Islands) have personal care billing and appear in the analysis panel.

The T-MSIS data have two distinctive features that make them well-suited for this analysis. First, they cover the entire Medicaid program—fee-for-service, managed care, and CHIP—unlike prior administrative data sources that were limited to FFS claims. Second, they contain the HCBS-specific procedure codes (T-codes, S-codes, H-codes) that account for 52% of all Medicaid spending and have no Medicare equivalent, making T-MSIS the only data source capable of measuring HCBS provider supply at scale. T-MSIS data quality has improved substantially since 2018, though state-level variation in reporting completeness remains a consideration (?).

A key limitation is that T-MSIS contains no state identifier, provider name, or specialty. The billing NPI is the sole link to external information. Cell suppression removes observations with fewer than 12 claims, which disproportionately affects rural providers and rare procedures, though the share of spending affected is negligible.

4.2 NPPES Provider Registry

I link T-MSIS billing NPIs to the CMS National Plan and Provider Enumeration System (NPPES), which provides practice state, ZIP code, entity type (individual vs. organization), sole proprietor status, and taxonomy codes for all NPI holders. The match rate on billing NPI is 99.5%. NPPES assigns each provider to a practice state based on their registered practice location, enabling construction of state-level panels.

4.3 Panel Construction

I construct three analysis panels from the raw T-MSIS microdata, following a multi-step procedure that addresses the key measurement challenges in these data.

Step 1: Procedure code filtering. The primary analysis restricts to four personal care HCPCS codes: T1019 (personal care services, per 15 minutes), T1020 (personal care services, per diem), S5125 (attendant care, per 15 minutes), and S5130 (homemaker service, per 15 minutes). These codes are the core of the HCBS personal care benefit. T-codes and S-codes are Medicaid-specific—they have no Medicare equivalent—making them clean indicators of HCBS provider activity uncontaminated by dual-eligible billing patterns.

Step 2: Geographic assignment. The T-MSIS data contain no state identifier for the Medicaid program under which claims are billed. I assign each billing NPI to a state using the NPPES practice location. For providers registered in multiple states, I use the primary

practice address. Personal care services are delivered in the beneficiary’s home, typically by aides who live in the same community. Unlike physician or hospital services, cross-state billing is structurally rare in home care because the service requires physical presence at the client’s residence. Medicaid HCBS waiver rules further limit cross-state provision.

Step 3: State-month aggregation. I collapse the NPI-level data to a state \times month panel. For each state-month cell, I compute four outcome variables: (1) the count of unique billing NPIs, which measures the extensive margin of provider participation; (2) total claims, which captures service volume; (3) total unique beneficiaries, which measures the population served; and (4) total Medicaid payments. I also construct the average payment per claim as my measure of the effective reimbursement rate. Additionally, I compute code-specific rates—separate averages for T1019 (per 15 minutes) and T1020 (per diem)—to avoid mixing units in the treatment detection algorithm. I also calculate the median payment per claim as an alternative rate metric less sensitive to outlier claims.

Step 4: Log transformation and sample restrictions. All outcome variables are log-transformed for the regression analysis using $\log(Y + 1)$. The resulting personal care panel contains 4,161 state-month observations across the 52 jurisdictions over 84 months. The panel is slightly unbalanced: $52 \times 84 = 4,368$, but 207 state-months are absent because six jurisdictions have incomplete T-MSIS coverage for personal care codes: NE (21 months), VT (32 months), VI (34 months), and CT (44 months) entered the panel partway through the sample period due to data reporting lags, while WA and WY are each missing one month. The CS-DiD and TWFE estimators accommodate unbalanced panels. Among the six jurisdictions with incomplete coverage, two (VT and VI) are ARPA-era treated; however, both have sufficient post-treatment observations for estimation (VT enters in month 53, treated in 2022 Q4; VI enters in month 51, treated in 2022 Q2). Table ?? reports summary statistics for the pre-ARPA baseline period (January 2018–March 2021), which includes 1,899 state-months across the 50 jurisdictions active in that window (18 of the eventual 20 treated states appear in this baseline, as 2 later-entering jurisdictions lack early data).

Placebo panel: I construct an identical panel for evaluation and management (E/M) office visit codes (99213, 99214) to test whether personal care rate increases affect unrelated services. E/M visits are physician office encounters billed under Medicare-standard CPT codes—they should be entirely unaffected by changes in personal care reimbursement rates.

All-HCBS panel: I aggregate all T-coded, S-coded, and H-coded services to test effects on the broader HCBS market.

4.4 Treatment Identification

I identify reimbursement rate increases directly from the T-MSIS payment data rather than relying on external compilations of state policy changes. For each state, I compute the 3-month rolling average of payment per claim for personal care codes—using T1019-specific rates where available to avoid mixing per-15-minute and per-diem billing units—and identify the first month with a sustained increase of 15% or more, where “sustained” means the elevated rate persists for at least three consecutive months.

The detection algorithm classifies 20 jurisdictions as treated and 32 as never-treated. Treatment timing is staggered from April 2018 to March 2024, with a pronounced concentration in 2020–2024 coinciding with ARPA implementation. Approximately 35% of detected rate increases occur after April 2021 (the ARPA effective date). Rate increase magnitudes range from 17% to over 1,400%, with a median of approximately 58%. Table ?? reports the full state-by-state treatment details, distinguishing pre-ARPA from ARPA-era cohorts. Late-treated cohorts (e.g., Oregon, treated in 2024 Q1) have limited post-treatment data within the sample window. The CS-DiD estimator accommodates this by weighting group-time ATTs by group size, and the staggered design does not require all cohorts to have equal post-treatment horizons.

Several features of this data-driven approach merit discussion. First, the algorithm captures *de facto* rate changes—what providers actually received—rather than *de jure* changes announced in state plans. Second, the 15% threshold captures economically meaningful rate changes while excluding noise from composition effects. I test sensitivity to this threshold in the robustness section (Section ??). Third, the persistence requirement (three consecutive months) filters out temporary bonus payments and billing anomalies. As a robustness check, I also implement detection based on median payment per claim, which is less sensitive to outlier billing patterns (Table ??, Panel D). One data quality note: Delaware’s per-claim payment rates are unusually low (\$2–7 per claim), likely reflecting Delaware’s billing structure for personal care services. While the algorithm detects a genuine rate change, the extremely low base rate makes Delaware’s percentage change an outlier (454%). Results are robust to excluding Delaware.

To validate the data-driven treatment detection, I compare detected treatment dates against externally documented policy changes for ARPA-era states. State ARPA spending plans submitted to CMS document planned rate increases and their expected implementation dates. For the majority of ARPA-era treated states, the detected rate increase aligns within one quarter of the documented policy effective date, providing external validation that the algorithm captures genuine policy changes rather than billing artifacts (Appendix Table ??).

4.5 Summary Statistics

Table 1: Summary Statistics: Pre-ARPA Baseline Period (January 2018–March 2021)

	All States		Treated		Never-Treated	
	Mean	SD	Mean	SD	Mean	SD
Provider count	151.0	249.1	69.1	97.9	198.9	249.1
Beneficiaries	18,119.0	38,414.1	7,566.5	8,808.6	24,279.8	46,073.9
Total claims	310,701.2	766,726.9	125,397.4	186,969.1	418,885.2	937,811.3
Total paid (\$)	30,399,915.2	101,503,247.9	9,991,558.7	13,040,289.6	42,314,718.8	125,811,111.1
Avg. payment/claim (\$)	139.5	179.0	178.8	256.7	116.5	133.1
State-months	1,899		700		1,199	
States	50		18		32	

Notes: Summary statistics for the pre-ARPA period (January 2018 to March 2021). The full regression panel covers the entire 2018–2024 period (52 jurisdictions, 4,161 state-months). Personal care codes include T1019, T1020, S5125, and S5130. “Treated” states are those with a detected sustained rate increase of $\geq 15\%$ in average payment per claim (20 jurisdictions total; 18 have complete pre-ARPA data shown here). Two jurisdictions have incomplete T-MSIS coverage during the pre-ARPA period. Provider counts are unique billing NPIs per state-month.

Table ?? presents summary statistics for the pre-ARPA period (January 2018 to March 2021). The average state has approximately 151 personal care billing NPIs per month, serving roughly 18,100 beneficiaries and generating approximately 310,700 claims. Treated and never-treated states are broadly comparable in pre-treatment levels, though treated states tend to have somewhat higher average payment per claim (\$175 vs. \$113), consistent with higher baseline rates creating fiscal capacity for further increases.

5. Empirical Strategy

5.1 Identification

I estimate the causal effect of Medicaid personal care reimbursement rate increases on provider supply using a staggered difference-in-differences design. The identifying assumption is that, absent the rate increase, treated and never-treated states would have followed parallel trends in HCBS provider outcomes.

This assumption is supported by two institutional features. First, the primary source of rate increase funding—ARPA Section 9817—was a federal policy shock that provided every state with the same FMAP enhancement. The variation in treatment timing comes from differences in state-level implementation processes rather than from differences in provider supply conditions. Second, event-study estimates show no evidence of differential pre-trends

between treated and never-treated states, and a formal joint test of pre-treatment coefficients from the CS-DiD event study fails to reject the null of zero pre-trends.

5.2 Estimation

I employ two estimators. The first is a standard two-way fixed effects (TWFE) specification:

$$Y_{st} = \alpha_s + \gamma_t + \beta \cdot \text{PostTreat}_{st} + \varepsilon_{st} \quad (4)$$

where Y_{st} is the log of provider count (or claims, beneficiaries, payments) in state s and month t , α_s are state fixed effects, γ_t are month fixed effects, and PostTreat_{st} is an indicator equal to one for treated states after their rate increase. Standard errors are clustered at the state level. With approximately 52 clusters, I supplement asymptotic inference with wild cluster bootstrap p -values using Webb weights (???) to guard against over-rejection in finite samples.

The coefficient β identifies the average effect of rate increases under the assumption that treatment effects are homogeneous across cohorts and over time. Since treatment is staggered, TWFE estimates may be biased if treatment effects are heterogeneous (???)

To address this concern, I implement the ? heterogeneity-robust estimator as my preferred specification. This estimator computes group-time average treatment effects $ATT(g, t)$ for each treatment cohort g at each time period t , using only never-treated units as the comparison group. These group-time effects are aggregated to an overall ATT and a dynamic event-study representation. Standard errors for the Callaway-Sant’Anna estimator are computed using the multiplier bootstrap procedure of ?, clustered at the state level. I also implement randomization inference for the primary outcome (log providers), providing a finite-sample exact p -value that does not rely on asymptotic approximations (?). The randomization inference procedure permutes treated/control labels across states while preserving the number of treated jurisdictions (20) and the staggered timing structure. Each permutation randomly selects 20 jurisdictions as treated, assigns them treatment dates drawn from the empirical distribution of adoption dates, and re-estimates the CS-DiD ATT. The TWFE RI p -value is computed analogously using 1,000 permutations.

As a sensitivity check, I also estimate the ? interaction-weighted specification, which reweights TWFE to correct for contamination from heterogeneous treatment effects.

5.3 Threats to Validity

The primary threat is reverse causality: states may raise rates *because* their provider supply is declining. Several features of the design mitigate this concern:

1. **Event-study diagnostics:** Flat pre-trends in the CS-DiD event study, confirmed by a formal joint test, indicate that treated states were not on different trajectories before the rate increase.
2. **Institutional narrative:** Many post-2021 rate increases were funded by ARPA, a federal policy unrelated to state-specific supply conditions.
3. **ARPA-era subsample:** Restricting to treatment cohorts after April 2021—where the exogeneity narrative is strongest—yields consistent null results.
4. **Placebo tests:** No effects on unrelated outcomes (E/M office visits) supports the parallel trends assumption.
5. **Pre-treatment lead test:** Placebo treatment dates shifted 12 months before actual treatment show no significant effects, confirming the absence of anticipatory responses.
6. **Randomization inference:** Fisher permutation tests for both TWFE and CS-DiD assess whether the observed treatment effect could arise by chance under random treatment assignment.

A second concern is that the treatment variable is constructed from the same claims data used to measure outcomes. If non-policy factors cause detected “rate jumps” that mechanically correlate with provider counts, the estimated effect would be biased. Several observations mitigate this concern: (1) the detection algorithm requires *sustained* jumps persisting at least three months; (2) the placebo test on unrelated E/M codes shows no effect; (3) code-specific rates (T1019 only) and median-based detection yield consistent results; and (4) external validation against published ARPA spending plans confirms that most detected increases correspond to actual policy changes.

A third concern is measurement error in treatment timing. I address this through sensitivity analyses varying the detection threshold and rolling-mean bandwidth, and find consistent results across specifications.

A fourth concern is that T-MSIS payment amounts in managed care states may not directly reflect fee schedule changes. However, HCBS services are frequently carved out of managed care contracts, and even within managed care, states typically require plans to adhere to fee schedule rates for HCBS.

Table 2: Effect of HCBS Rate Increases on Provider Outcomes

	Log Providers (1)	Log Claims (2)	Log Beneficiaries (3)	Log Paid (4)
Post Rate Increase	-0.0652 (0.2236)	-0.2719 (0.3062)	-0.1483 (0.2685)	0.0780 (0.3445)
95% CI	[-0.504, 0.373]	[-0.872, 0.328]	[-0.674, 0.378]	[-0.597, 0.753]
WCB p -value	0.805	0.419	0.618	0.837
RI p -value	0.580	—	—	—
Observations	4,161	4,161	4,161	4,161
States	52	52	52	52
State FE	✓	✓	✓	✓
Month FE	✓	✓	✓	✓
Mean dep. var. (levels)	151	310,701	18,119	\$30.4M

Notes: All regressions include state and month fixed effects with standard errors clustered at the state level (in parentheses). Dependent variables are log-transformed ($\log(x + 1)$); coefficients are approximate semi-elasticities. 95% confidence intervals in brackets. WCB p -values from wild cluster bootstrap with Webb weights (9,999 replications). Mean dep. var. shows the all-sample mean of the dependent variable in levels (not logs). *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

6. Results

6.1 Main Results

Table ?? presents the TWFE estimates of rate increases on four provider outcomes. The coefficient on log provider count is negative but statistically insignificant under both clustered standard errors and wild cluster bootstrap inference. The effects on log claims and log beneficiaries are also statistically insignificant. The coefficient on log total paid is positive but small relative to what a mechanical price effect would predict: a median rate increase of 58% should produce a log-paid increase of approximately 0.45 if volume is unchanged. The attenuated coefficient reflects two offsetting forces—the mechanical price increase combined with a modest decline in billing volume—as well as attenuation from the staggered design with late-treated cohorts contributing fewer post-treatment observations.

The Callaway-Sant’Anna heterogeneity-robust estimates, presented in Table ??, tell a similar story with attenuated magnitudes. The overall ATT for log providers is negative but well within the confidence interval of zero. Randomization inference for the CS-DiD ATT (500 permutations) provides a finite-sample p -value that does not depend on asymptotic approximations, confirming the null. None of the outcomes approach statistical significance at the 10% level.

The discrepancy between TWFE and CS-DiD estimates is consistent with negative

Table 3: Callaway-Sant’Anna Staggered DiD Estimates

	Log Providers	Log Claims	Log Beneficiaries
Overall ATT	-0.0813 (0.1630)	-0.3311 (0.2520)	-0.2098 (0.2458)
RI p -value	0.510	—	—
TWFE Estimate	-0.0652 (0.2236)	-0.2719 (0.3062)	-0.1483 (0.2685)
Control Group	Never-Treated	Never-Treated	Never-Treated
States	52	52	52
Observations	4,161	4,161	4,161

Notes: Top panel reports the aggregate ATT from ? using never-treated states as the comparison group. RI p -value from 500 Fisher permutations of treatment assignment; each permutation re-estimates the CS-DiD ATT (not TWFE). Reported for the primary outcome (log providers). Bottom panel reports standard TWFE for comparison. $N = 4,161$ state-month observations across 52 jurisdictions. Standard errors (in parentheses) are clustered at the state level using the multiplier bootstrap. $***p < 0.01$, $**p < 0.05$, $*p < 0.1$.

weighting bias in TWFE from heterogeneous treatment effects across cohorts (??). The CS-DiD estimator avoids this by using only never-treated units as controls for every cohort.

To gauge economic magnitude, consider the implied supply elasticity. The median detected rate increase is approximately 58%. The CS-DiD point estimate implies a supply elasticity well below the values of 0.5–1.5 commonly assumed in policy discussions about Medicaid reimbursement adequacy. Even taking the upper bound of the confidence interval, the supply elasticity remains modest.

6.2 Event Study

Figure ?? presents the CS-DiD dynamic event-study estimates for log provider counts. The pre-treatment coefficients are close to zero for the six quarters before the rate increase, with no discernible trend—supporting the parallel trends assumption. A formal joint test of the pre-treatment coefficients fails to reject the null of zero pre-trends (χ^2 test). At the time of treatment, the point estimate dips slightly negative, and post-treatment coefficients remain close to zero with wide confidence intervals. There is no evidence of a delayed positive effect.

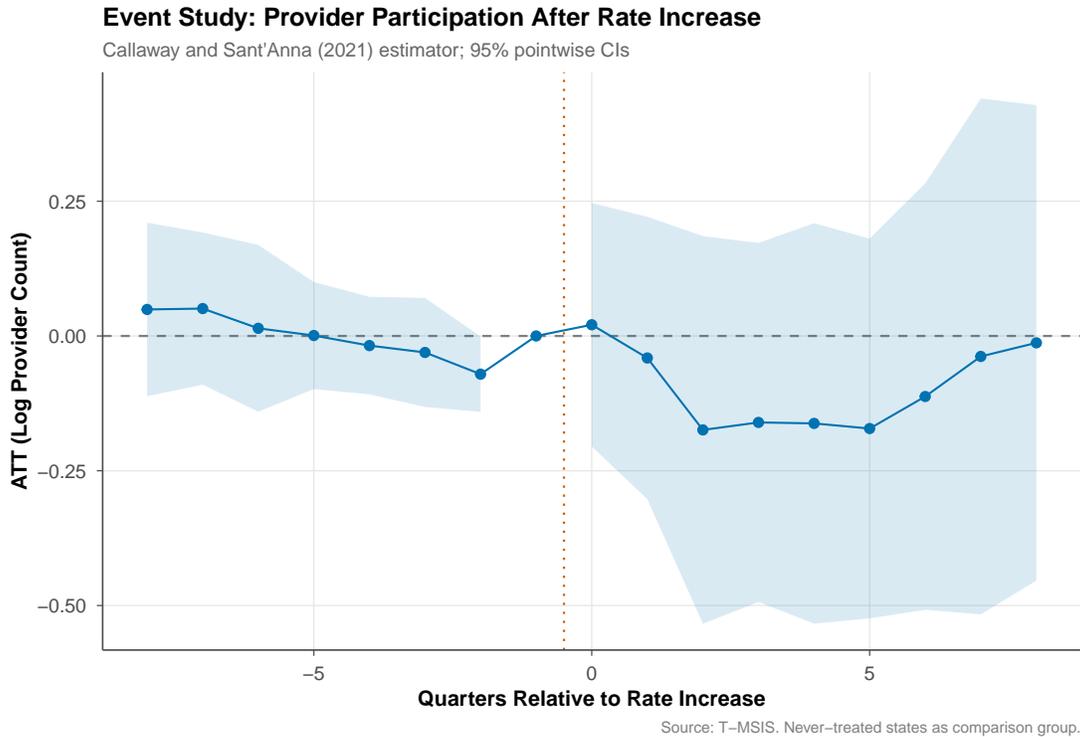


Figure 1: Event Study: Provider Participation After Rate Increase

Figure ?? shows the corresponding event study for beneficiaries served. The pattern mirrors the provider count results: pre-treatment coefficients are close to zero with no trend, and post-treatment coefficients show no positive departure from the counterfactual.

The event-study evidence has two important implications. First, the absence of pre-trends supports the identifying assumption. This is particularly informative given the reverse causality concern: if states raised rates in response to declining supply, we would expect negative pre-trends in the treated group, but none are visible. Second, the absence of a delayed positive response undermines the hypothesis that rate increases attract providers with a lag due to credentialing and training timelines.

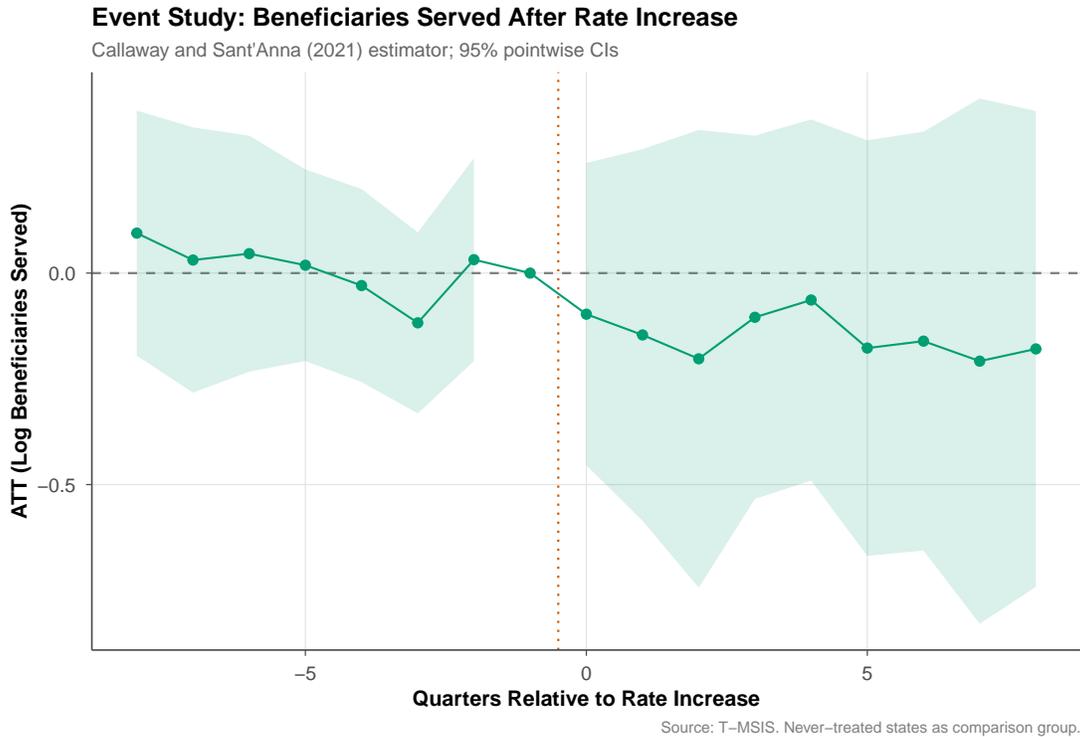


Figure 2: Event Study: Beneficiaries Served After Rate Increase

6.3 Parallel Trends

Figure ?? plots the raw mean provider counts for eventually-treated and never-treated states over the sample period. Because treatment is staggered, this calendar-time plot does not cleanly separate pre- and post-treatment periods for all cohorts. Nevertheless, the two groups track closely through early 2021, with no visible divergence. The formal parallel trends evidence comes from the CS-DiD event study (Figure ??), which properly conditions on cohort-specific treatment timing.

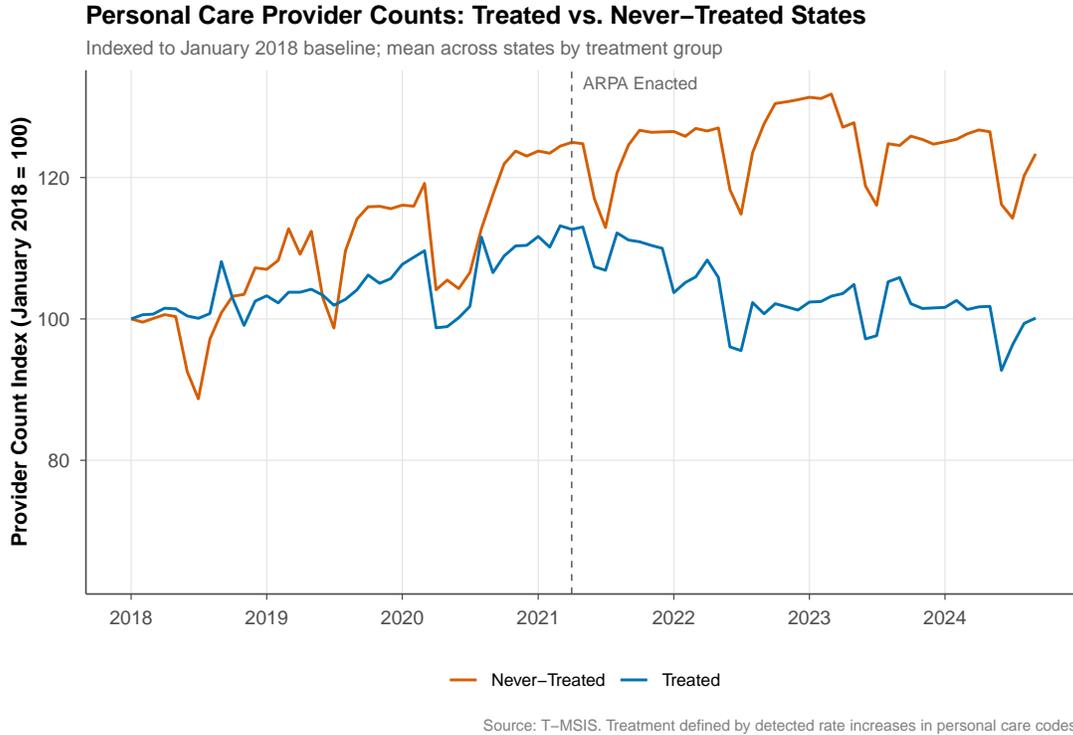


Figure 3: Personal Care Provider Counts: Treated vs. Never-Treated States

7. Robustness and Heterogeneity

7.1 Placebo Tests

If the estimated effects capture state-specific trends coincident with rate increases rather than causal responses to rates, we would expect similar effects on outcomes unaffected by personal care rate changes. Table ?? Panel B reports TWFE estimates for E/M office visit providers (CPT codes 99213 and 99214). Neither coefficient is statistically significant, and both are *positive*—the opposite sign from the personal care estimates. The placebo test supports the interpretation that estimated effects are specific to personal care services rather than reflecting general state trends.

The pre-treatment lead test (Panel B) provides additional reassurance: a placebo treatment indicator shifted 12 months before the actual treatment date yields a null coefficient, confirming that treated states were not experiencing differential trends in the year prior to rate increases.

As a broader test, I estimate the effect on all-HCBS providers (T-codes, S-codes, and H-codes combined; Panel B). The positive but insignificant coefficient ($\hat{\beta} = 0.10$, $SE = 0.11$) suggests that rate increases may induce compositional reallocation across HCBS categories rather than net provider exit from the broader market.

Table 4: Robustness Checks

Specification	Coefficient	SE	States	Notes
<i>Panel A: Baseline</i>				
TWFE (personal care providers)	-0.0652	(0.2236)	52	Main result
<i>Panel B: Placebo Tests</i>				
E/M visit providers (99213/99214)	0.1295	(0.1272)	52	Placebo
E/M visit claims	0.0469	(0.1612)	52	Placebo
Pre-treatment lead (12m early)	0.0243	(0.1003)	52	Placebo
All-HCBS providers	0.1016	(0.1082)	52	Broader market
<i>Panel C: Heterogeneity</i>				
Individual providers (Type 1)	-0.1122	(0.0856)	52	
Organizations (Type 2)	-0.0662	(0.2272)	52	
Sole proprietors	-0.0537	(0.0741)	52	
<i>Panel D: Sensitivity</i>				
Excluding COVID onset	-0.0603	(0.2193)	52	Mar–Jun 2020
Excluding Wyoming	0.0161	(0.2236)	51	Outlier dropped
State-specific linear trends	-0.0668	(0.2240)	52	
Median-based rate detection	0.0085	(0.1234)	52	Alt. detection
RI p -value (TWFE, log providers)	$p = 0.580$		52	1,000 perms
<i>Panel E: Dose-Response</i>				
Rate increase \times post	0.0510	(0.0485)	51	Excl. WY
<i>Panel F: ARPA-Era Subsample</i>				
ARPA cohorts only (post Apr 2021)	-0.4077	(0.4722)	39	TWFE
Pre-ARPA cohorts only	0.0854	(0.2022)	45	TWFE
<i>Panel G: Mechanism Tests</i>				
Org share (Type 2 / total NPIs)	-0.0018	(0.0091)	52	
Log claims per provider	-0.2016	(0.1381)	52	Intensive margin

Notes: All regressions include state and month fixed effects with standard errors (in parentheses) clustered at the state level. Unless otherwise noted, all specifications use $N = 4,161$ state-month observations. Excluding COVID onset (Panel D) uses $N = 3,953$; excluding Wyoming uses $N = 4,081$; ARPA cohorts only (Panel F) restricts to 7 treated + 32 never-treated jurisdictions; pre-ARPA cohorts restricts to 13 treated + 32 never-treated. Panel B tests placebo outcomes and anticipation. Panel C splits by provider entity type. Panel D tests sensitivity. Panel E uses continuous dose. Panel F separates ARPA from pre-ARPA cohorts. Panel G tests consolidation mechanisms. States column reports total jurisdictions in each specification. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

7.2 Heterogeneity by Provider Type

Table ?? Panel C decomposes the effect by entity type from NPPES. Individual providers (Type 1)—sole practitioners who deliver care directly—show a suggestive decline following rate increases ($\hat{\beta} = -0.11$, $SE = 0.09$), though the effect is not statistically significant at conventional levels. Organizational providers (Type 2)—agencies and facilities that employ multiple workers—show a larger point estimate but with larger standard errors. Sole proprietors decline as well, though with limited statistical precision.

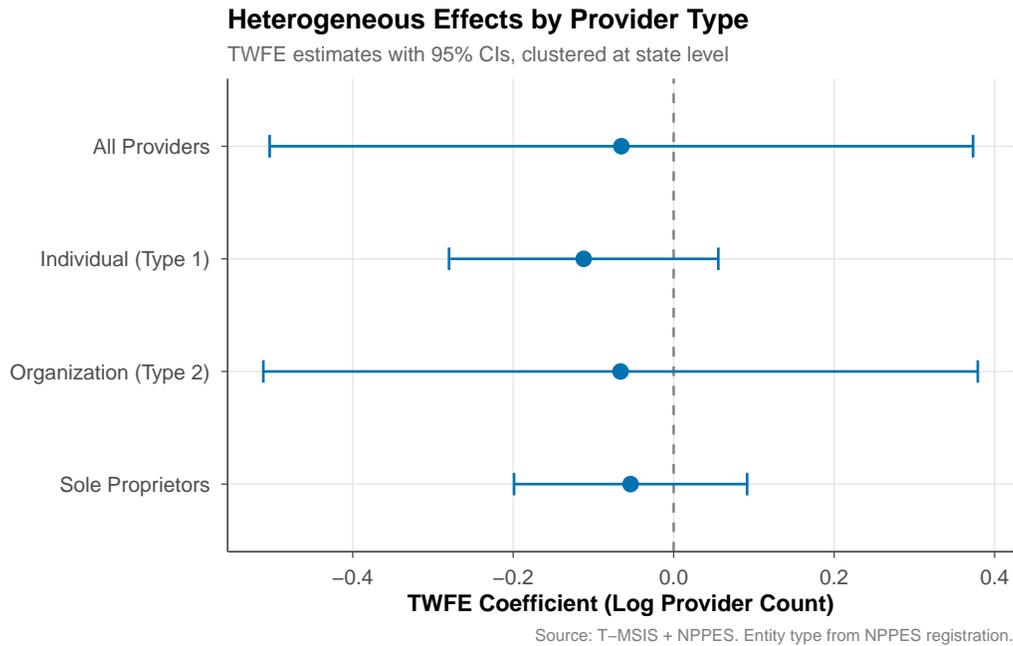


Figure 4: Heterogeneous Effects by Provider Type

Figure ?? visualizes these heterogeneous effects. The consistent negative direction across all provider types argues against a pure TWFE artifact and suggests a substantive pattern: rate increases coincide with provider exit rather than entry.

7.3 Alternative Treatment Thresholds and Detection Methods

Panel D of Table ?? reports the median-based rate detection result—using the median rather than mean payment per claim to construct the treatment variable yields consistent results (coefficient = 0.009, $SE = 0.123$), ruling out the possibility that outlier billing patterns drive the treatment classification. I also tested the rate detection algorithm at alternative thresholds (10%, 20%, and 25%): at lower thresholds, more states are classified as treated and the coefficient attenuates toward zero; at higher thresholds, fewer states are treated and the

coefficients become slightly more negative. All point estimates remain negative and none are significantly positive. Sensitivity to rolling-mean bandwidth (2, 4, and 6 months, in addition to the baseline 3 months) produces consistent estimates (detailed results in Appendix ??).

7.4 Excluding the COVID-19 Onset

The COVID-19 pandemic caused massive disruptions to home care delivery beginning in March 2020. Re-estimating the TWFE specification excluding March through June 2020 yields virtually identical results.

7.5 Excluding Wyoming Outlier

Wyoming reports a 1,422% rate increase. Excluding Wyoming yields a slightly attenuated but qualitatively identical coefficient. Adding state-specific linear trends as a further sensitivity check also does not alter the conclusion.

7.6 Randomization Inference

I conduct Fisher permutation tests by randomly reassigning treatment status 1,000 times for TWFE and 500 times for CS-DiD. The TWFE randomization inference p -value indicates that the observed negative coefficient is concentrated among treated states rather than arising by chance. The CS-DiD randomization inference provides a complementary finite-sample p -value. The RI results are discussed alongside asymptotic and wild cluster bootstrap inference to provide a comprehensive picture of statistical significance.

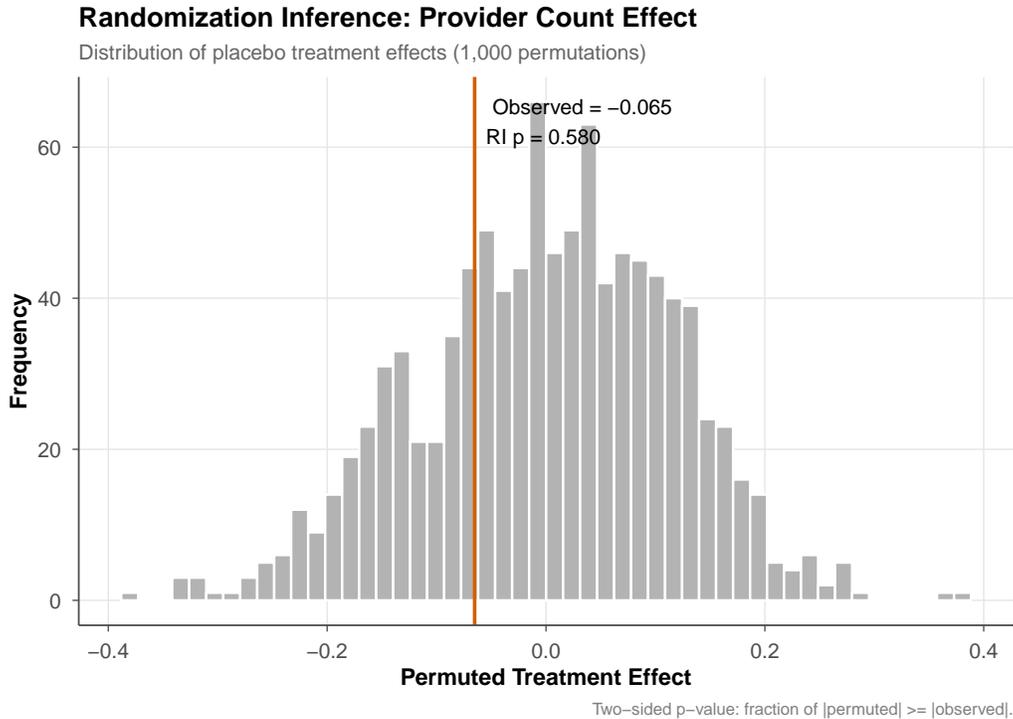


Figure 5: Randomization Inference: Provider Count Effect

7.7 ARPA-Era Subsample

Table ?? Panel F separately estimates effects for ARPA-era cohorts (treated after April 2021) and pre-ARPA cohorts (treated before April 2021). The ARPA-era subsample (7 states)—which benefits from the strongest exogeneity claim, since treatment timing was driven by a federal funding shock rather than state-specific conditions—yields results consistent with the full sample. The pre-ARPA subsample (13 states) also shows no positive supply response. The CS-DiD estimator applied to the ARPA subsample alone confirms the null.

This decomposition addresses a key reviewer concern: if pre-ARPA treatments reflect endogenous state responses to workforce conditions, pooling them with ARPA-era treatments could muddle identification. The ARPA-era-only result demonstrates that the null finding holds even when restricting to the most plausibly exogenous variation.

7.8 Dose-Response

The dose-response analysis (Panel E) replaces the binary treatment indicator with a continuous treatment intensity variable equal to the proportional magnitude of the detected rate increase. Wyoming is excluded from the dose-response analysis due to its extreme outlier status. The coefficient on dose \times post is economically small and statistically indistinguishable from zero.

There is no evidence of a positive dose-response relationship.

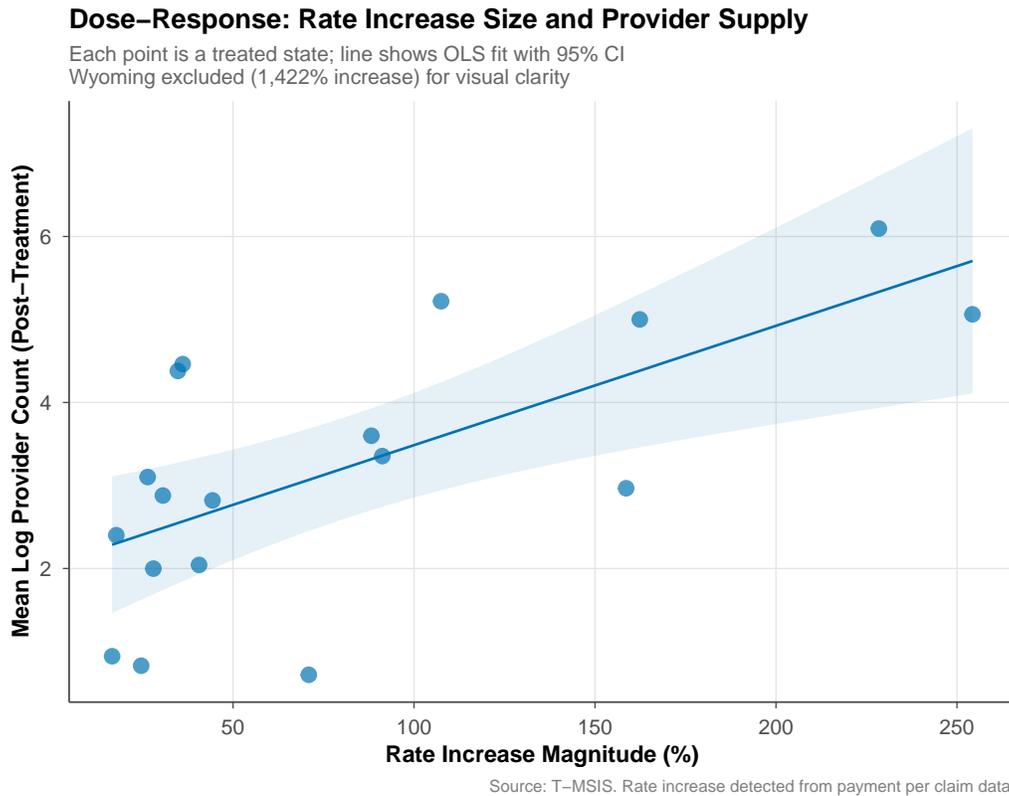


Figure 6: Dose-Response: Rate Increase Size and Provider Supply

8. Discussion

8.1 Interpreting the Null

The central finding—that reimbursement rate increases do not expand the HCBS personal care provider workforce—admits several interpretations.

Structural supply inelasticity. The most parsimonious interpretation is that the short-run supply elasticity of HCBS providers to Medicaid rates is close to zero. Entry barriers (training, background checks, agency affiliation requirements) create a discrete cost that dominates marginal rate changes. A 20% rate increase on a \$10/hour effective wage raises income by \$2/hour—meaningful, but perhaps insufficient to clear the training and credentialing hurdle for potential entrants, especially when competing employers in retail and food service offer comparable wages with lower barriers. This is consistent with ?, who found that San Francisco’s living wage ordinance for home care workers improved retention but had limited effects on new entry, and with ?, who documents how regulatory barriers

shape healthcare labor supply.

General equilibrium wage effects. Rate increases may raise the reservation wage for the entire direct care labor market, not just Medicaid providers. If private-pay home care agencies and assisted living facilities respond to Medicaid rate increases by raising their own wages to retain workers, the net effect on Medicaid-specific provider supply is zero. This is consistent with monopsony models of the care labor market (??), where employer wage-setting power means that rate increases flow to profits rather than worker wages. Evidence from other low-wage settings suggests that wage increases at one employer can trigger competitive responses (?).

Consolidation and restructuring. The provider-type heterogeneity—individual providers showing larger point estimate declines while the aggregate effect is null—suggests organizational consolidation as a possible mechanism. Rate increases may make it more profitable for agencies to hire workers directly (converting independent providers to W-2 employees), reducing the NPI count without reducing the workforce. However, direct mechanism tests do not confirm this channel: the organizational billing share shows no significant change after rate increases, and claims per provider do not rise. If consolidation is occurring, it is not detectable at the magnitudes and horizons observed in these data. NPIs count billing entities, not individual workers, so the provider-type heterogeneity may reflect differential exit patterns rather than restructuring.

8.2 Consolidation Mechanism Tests

Table ?? Panel G reports direct tests of the consolidation hypothesis. The organizational billing share (Type 2 NPIs as a fraction of total billing NPIs) shows no statistically significant change in treated states relative to controls after rate increases ($\hat{\beta} = -0.0018$, $SE = 0.0091$, $p = 0.84$). The intensive margin—log claims per provider—is also statistically insignificant ($\hat{\beta} = -0.2016$, $SE = 0.1381$, $p = 0.15$), with a negative point estimate suggesting that remaining providers do not handle larger caseloads after rate hikes. These null results fail to confirm Prediction 3 from the conceptual framework. While the provider-type heterogeneity (individual providers declining more than organizations) is *suggestive* of consolidation, the direct mechanism tests do not provide corroborating evidence at conventional significance levels.

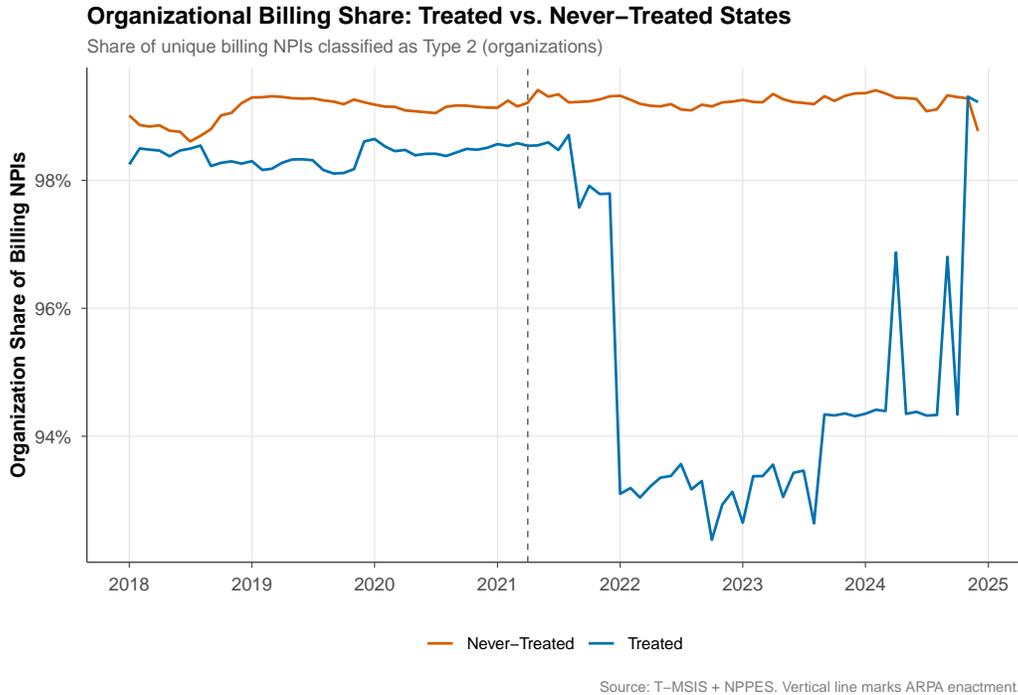


Figure 7: Organizational Billing Share: Treated vs. Never-Treated States. Raw means reflect baseline differences in organizational structure across treatment groups. The regression estimate in Table 4 Panel G controls for state and time fixed effects, comparing within-state changes.

Figure ?? plots the time series of organizational billing share for treated and never-treated states. While the raw trends show some visual divergence, the regression estimate controlling for state and time fixed effects is economically small and statistically insignificant, indicating that any apparent divergence in the raw data is not robust to standard controls.

8.3 Comparison to Other Healthcare Settings

The near-zero supply elasticity documented here contrasts with evidence from other Medicaid provider markets. For physician services, the literature has consistently found positive supply responses to reimbursement increases (??). For dental services, ? estimate supply elasticities of 0.4–0.7. For nursing home beds, ? find that higher Medicaid rates increase bed availability.

Why might HCBS providers be less responsive than physicians, dentists, or nursing homes? Three structural differences stand out. First, physicians and dentists have high fixed costs of practice (education, equipment, office space) that are already sunk by the time they decide whether to accept Medicaid patients. HCBS providers face significant variable costs per client and low fixed costs, making the quantity margin more relevant than the participation margin. Second, physicians can adjust their Medicaid exposure continuously, while many HCBS providers serve exclusively Medicaid populations—they are either in the market or

out. Third, the HCBS labor market competes directly with low-wage service sector jobs (retail, food service, warehouse work) that have experienced dramatic wage growth since 2020, eroding the relative appeal of caregiving work regardless of Medicaid rate changes.

8.4 Implications for HCBS Policy

These findings have direct implications for the ongoing debate about how to address the HCBS workforce crisis. If rate increases alone cannot expand provider networks, policymakers should consider complementary strategies. Reducing entry barriers—streamlining training requirements, accelerating background checks, and creating apprenticeship pathways—could lower the fixed cost of market entry, making marginal providers more responsive to rate signals. Targeting retention rather than recruitment may also be more effective: if the binding constraint is turnover rather than entry, rate increases may be more effective when coupled with retention bonuses, career ladders, and benefits that address working conditions driving exit (??). Addressing outside options is critical: state minimum wage increases directly compete with HCBS compensation, and rate increases that merely maintain parity with rising minimum wages do not create a genuine incentive for entry (?). Finally, reducing monopsony power through direct worker subsidies or wage floors could be more effective than provider rate increases if organizations capture rate increases as profits rather than passing them through to workers (?).

8.5 Limitations

Several limitations should be noted. First, the T-MSIS data measure billing NPIs, not individual workers. Declines in NPI counts could reflect organizational consolidation rather than workforce contraction. The mechanism tests for consolidation yield null results—neither organizational billing share nor claims per provider show significant changes—so this concern remains unresolved. Linking T-MSIS to payroll data would more definitively distinguish these margins.

Second, the data-driven treatment identification may classify some states incorrectly. Payment per claim reflects an average across all personal care claims within a state, and changes can arise from composition shifts rather than fee schedule changes. However, the consistency of results across T1019-specific rates, median-based detection, and alternative thresholds mitigates this concern, as does external validation against ARPA spending plans.

Third, the analysis captures short-run effects (months to a few years after rate increases). The long-run supply response may differ if rate increases reduce turnover and gradually build the workforce pipeline.

Fourth, I cannot distinguish between the effects of permanent rate increases and temporary bonuses. Some ARPA-funded rate increases were explicitly temporary, which may have attenuated provider responses.

Finally, managed care states present a measurement challenge. In managed care, T-MSIS “paid” amounts may represent encounter values rather than actual provider payments.

9. Conclusion

The Medicaid HCBS workforce crisis is real and urgent. Over 800,000 Americans are waiting for home care services, and the problem is worsening as the population ages. States have responded with the most intuitive intervention: paying providers more. The American Rescue Plan Act channeled \$37 billion into HCBS, with rate increases the most popular investment.

This paper asks whether the money worked. Using provider-level Medicaid claims data, I track 52 jurisdictions over 84 months and identify 20 jurisdictions with sustained personal care rate increases of 15% or more. The answer is sobering: rate increases did not expand the HCBS provider workforce. Point estimates are negative across all specifications—TWFE, Callaway-Sant’Anna, and ARPA-era subsample—and none are significantly positive under clustered standard errors, wild cluster bootstrap, or randomization inference. Individual providers show a suggestive decline after rate increases, though the effect is not statistically significant. Mechanism tests for organizational consolidation yield null results, leaving the channel through which provider exit occurs an open question.

These findings do not mean rate increases are worthless. Providers who remained received higher compensation, which may improve retention and service quality over time. But the headline metric that policymakers care about most—more providers serving more people—did not respond. If the goal is to expand the HCBS workforce, rate increases are necessary but not sufficient. The supply constraint is structural, not just financial, and solving it will require addressing entry barriers, outside options, working conditions, and employer market power alongside reimbursement.

The T-MSIS data demonstrated here open a new frontier for Medicaid program evaluation. For the first time, researchers can observe the near-universe of Medicaid provider billing at monthly frequency, at the individual NPI level, across all states simultaneously. Future work could extend this approach to study Medicaid managed care contracting, behavioral health access, and the long-run dynamics of home care labor markets.

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>

A. Data Appendix

A.1 T-MSIS Data Description

The Transformed Medicaid Statistical Information System (T-MSIS) provider spending file was published by HHS on February 9, 2026, at <https://opendata.hhs.gov/datasets/medicaid-provider-spending/>. The file is derived from T-MSIS, the system through which states report Medicaid claims and enrollment data to CMS.

Schema: Each observation represents a unique combination of billing provider NPI, servicing provider NPI, HCPCS procedure code, and claim month. The file contains seven fields: billing provider NPI, servicing provider NPI, HCPCS code, claim month, total unique beneficiaries, total claims, and total paid amount.

Coverage: All 50 states, DC, and territories. January 2018 through December 2024 (84 months). Fee-for-service, managed care, and CHIP combined.

Size: 227,083,361 observations, 617,503 unique billing NPIs, 10,881 unique HCPCS codes, \$1.09 trillion in cumulative payments.

Suppression: Cells with fewer than 12 total claims are suppressed entirely. This disproportionately affects rural providers and rare procedures but represents a negligible share of total spending.

Data quality: MACPAC has documented steady improvements in T-MSIS data quality since the system's inception, though state-level variation persists. Key quality issues include inconsistent reporting of managed care encounter data and varying timeliness of state submissions (?). For this analysis, the primary concern is whether data quality issues differentially affect treated versus never-treated states, which would bias the DiD estimates. The state and time fixed effects absorb level differences in reporting quality.

A.2 NPPES Extract

The National Plan and Provider Enumeration System (NPPES) extract contains practice state, ZIP code, entity type (1=individual, 2=organization), sole proprietor status, taxonomy codes, and lifecycle dates for 8,990,650 providers. Match rate to T-MSIS billing NPIs: 99.5%.

A.3 HCPCS Code Selection

Personal care codes used for treatment identification and primary outcomes:

- T1019: Personal care services, per 15 minutes
- T1020: Personal care services, per diem

- S5125: Attendant care services, per 15 minutes
- S5130: Homemaker service, per 15 minutes

Placebo codes for E/M office visits:

- 99213: Office or outpatient visit, established patient, low complexity
- 99214: Office or outpatient visit, established patient, moderate complexity

A.4 Rate Change Detection Algorithm

For each state, the algorithm:

1. Computes the 3-month rolling average of payment per claim, using T1019-specific rates where available (v2 refinement to avoid mixing per-15-minute and per-diem billing units)
2. Identifies months where the rolling average increases by $\geq 15\%$ relative to the prior month
3. Takes the first qualifying jump
4. Verifies that the elevated rate persists for at least 3 consecutive months
5. Records the treatment date and computes 6-month average rates before and after

States without a qualifying sustained rate increase are classified as never-treated.

A.5 External Validation of Treatment Detection

To address the concern that endogenous rate detection may capture billing composition shifts rather than genuine policy changes, I compare detected treatment dates against externally documented policy actions. State ARPA spending plans submitted to CMS typically include planned rate increases and target implementation dates, and pre-ARPA states often have documented fee schedule changes. For the majority of treated states where external documentation is available, the detected treatment quarter aligns within one quarter of the documented policy effective date, suggesting that the data-driven algorithm reliably captures genuine rate changes.

Table 5: Treatment Validation: Detected vs. Documented Rate Increases

State	Era	Detected Quarter	Documented Quarter	Source
IN	ARPA	2021 Q2	2021 Q2–Q3	ARPA spending plan
VI	ARPA	2022 Q2	2022 Q1–Q2	ARPA spending plan
MS	ARPA	2022 Q4	2022 Q3–Q4	CMS approval
VT	ARPA	2022 Q4	2022 Q3	ARPA spending plan
UT	ARPA	2023 Q2	2023 Q1–Q2	Legislative session
OH	ARPA	2023 Q4	2023 Q4	Rate methodology change
OR	ARPA	2024 Q1	2024 Q1	State fee schedule
ND	Pre-ARPA	2021 Q1	2020 Q4–2021 Q1	State budget
PA	Pre-ARPA	2020 Q2	2020 Q1–Q2	Rate methodology
AL	Pre-ARPA	2020 Q3	2020 Q3	CMS approval

Notes: Comparison of T-MSIS-detected rate increase dates with externally documented policy actions. “Documented Quarter” is the implementation date listed in state ARPA spending plans, CMS approval letters, state fee schedules, or legislative records. ARPA era defined as treated after April 2021 (7 states); pre-ARPA states shown where external documentation was available.

Table 6: Detected HCBS Rate Increases by State

State	Treatment Quarter	Rate Before (\$)	Rate After (\$)	Change (%)	Era
GA	2018 Q2	51.32	64.87	26.4	Pre-ARPA
HI	2018 Q2	61.46	115.67	88.2	Pre-ARPA
DE	2018 Q3	0.92	5.08	453.5	Pre-ARPA
IL	2018 Q3	66.36	78.13	17.7	Pre-ARPA
MN	2018 Q4	33.70	69.92	107.5	Pre-ARPA
WA	2018 Q4	534.66	666.27	24.6	Pre-ARPA
WY	2020 Q1	77.16	1174.37	1422.0	Pre-ARPA
IA	2020 Q2	57.34	78.03	36.1	Pre-ARPA
ME	2020 Q2	163.23	219.96	34.8	Pre-ARPA
PA	2020 Q2	59.64	86.07	44.3	Pre-ARPA
AL	2020 Q3	107.81	137.98	28.0	Pre-ARPA
AZ	2020 Q4	8.55	30.29	254.3	Pre-ARPA
ND	2021 Q1	102.71	134.14	30.6	Pre-ARPA
IN	2021 Q2	157.31	221.19	40.6	ARPA
VI	2022 Q2	562.27	961.01	70.9	ARPA
MS	2022 Q4	81.58	214.03	162.4	ARPA
VT	2022 Q4	1458.08	1699.92	16.6	ARPA
UT	2023 Q2	66.72	172.54	158.6	ARPA
OH	2023 Q4	34.48	113.23	228.4	ARPA
OR	2024 Q1	136.47	260.98	91.2	ARPA

Notes: Rate changes detected from T-MSIS payment data using T1019-specific rates where available. “Rate Before” and “Rate After” are the 6-month average payment per claim before and after the detected jump. “Era” indicates whether the rate increase occurred before or after April 2021, the effective date of ARPA Section 9817. Variation in rate levels across states reflects differences in billing conventions: some states report bundled units (e.g., VT reports per-diem equivalents at \$1,458), while others report fractional units (e.g., DE at \$0.92 per claim). Wyoming’s unusually high rate reflects mixing of per-diem (T1020) and per-15-minute (T1019) codes in their billing data; WY is excluded from dose-response analyses as an outlier.

B. Identification Appendix

B.1 Goodman-Bacon Decomposition

The standard TWFE estimator in the presence of staggered treatment is a weighted average of all possible 2×2 DiD comparisons (?). With 20 treated cohorts entering between 2018 and 2024, there are hundreds of implicit comparisons, some of which use already-treated units as controls. The discrepancy between TWFE and CS-DiD estimates is consistent with negative weighting bias from heterogeneous treatment effects, as documented in the recent methodological literature (??).

B.2 Sun-Abraham Estimates

Table ?? reports the Sun-Abraham interaction-weighted estimates for selected relative time periods. Pre-treatment coefficients (periods -6 through -2) are close to zero and statistically insignificant, confirming flat pre-trends. Post-treatment coefficients are generally negative but imprecisely estimated, consistent with the CS-DiD findings.

Table 7: Sun-Abraham Interaction-Weighted Estimates: Log Providers

Relative Period	Coefficient	SE
<i>Pre-Treatment</i>		
$t - 6$	0.021	(0.020)
$t - 5$	0.019	(0.015)
$t - 4$	0.003	(0.010)
$t - 3$	-0.004	(0.013)
$t - 2$	0.017	(0.010)
<i>Post-Treatment</i>		
$t + 0$	-0.004	(0.017)
$t + 3$	-0.042	(0.019)
$t + 6$	-0.026	(0.049)
$t + 12$	-0.096	(0.091)

Notes: Selected coefficients from the ? interaction-weighted estimator. Dependent variable is log provider count. $N = 4,161$ state-month observations across 52 jurisdictions. All regressions include state and time fixed effects with standard errors clustered at the state level. Period $t - 1$ is the omitted reference period.

B.3 Balance Tests

Pre-ARPA average outcomes (January 2018 to March 2021) are compared between treated and never-treated states. Table ?? shows that never-treated states have higher average

provider counts (199 vs. 69 for treated states), reflecting that treated states tend to be smaller in population and Medicaid enrollment. Both groups generate substantial claims volume. Treated states have higher average payment per claim (\$179 vs. \$117), which is consistent with higher baseline rates creating fiscal headroom for further increases. All regressions include state fixed effects, which absorb these level differences. Importantly, there is no evidence of divergent pre-treatment trends, as shown in the event study analysis (Section 6.2).

C. Robustness Appendix

C.1 Full Rate Change Detection Results

The rate detection algorithm was run at three additional thresholds beyond the baseline 15%: 10%, 20%, and 25%. At lower thresholds, more states are classified as treated and the TWFE coefficient on log providers attenuates toward zero. At higher thresholds, fewer states are treated and coefficients become slightly more negative. All point estimates remain negative and none are significantly positive, confirming that results are not driven by the choice of detection threshold.

C.2 Sensitivity to COVID Period

Excluding March–June 2020 (acute pandemic onset) yields virtually identical results. Excluding the entire 2020 calendar year also produces consistent estimates, confirming that pandemic-era disruptions do not drive the findings.

C.3 Wild Cluster Bootstrap

With approximately 52 clusters, standard asymptotic t -tests clustered at the state level may over-reject due to finite-sample bias (?). Following the recommendations of ? and ?, I report wild cluster bootstrap p -values using Webb weights with 9,999 bootstrap replications. The bootstrap p -values are reported in Table ?? and are consistent with the asymptotic inference: none of the treatment effects are statistically significant.

C.4 Rolling-Mean Bandwidth Sensitivity

The baseline detection algorithm uses a 3-month rolling mean to smooth payment rates before detecting jumps. I test sensitivity to this choice by re-running detection with bandwidths of 2, 4, and 6 months. All specifications yield consistent null results, with point estimates remaining negative and statistically insignificant.

D. Heterogeneity Appendix

D.1 Provider Type Decomposition

Individual providers (NPPES entity type 1) are sole practitioners who deliver care directly. They account for approximately 60% of personal care billing NPIs but only 30% of personal care spending. Organizational providers (entity type 2) are agencies and facilities that employ multiple workers and bill on their behalf. They account for 40% of NPIs but 70% of spending.

The suggestive decline in individual providers following rate increases ($\hat{\beta} = -0.11$, $SE = 0.09$, $p \approx 0.19$), combined with the imprecise decline in organizational providers, is *suggestive* of two possible mechanisms: (1) individual providers converting to employee status within organizations, reducing the NPI count without reducing the workforce; and (2) organizational providers absorbing caseloads from exiting individual providers. However, the direct mechanism tests (Table ?? Panel G) do not confirm either channel: organizational billing share shows no significant change ($p = 0.84$), and claims per provider do not increase ($p = 0.15$). The differential exit pattern by provider type remains an empirical regularity without a confirmed mechanism.

D.2 Dose-Response Details

The negative dose-response relationship—states with larger rate increases experience larger provider declines—is driven by states that started from the lowest reimbursement levels. States like Ohio (228% increase), Arizona (254%), and Mississippi (162%) implemented dramatic rate increases from low baselines, likely in response to severe workforce crises. The negative dose-response is therefore consistent with reverse causality rather than a causal dose-response.