

Minimum Wages and Medicaid’s Invisible Workforce: Provider Supply Responses to Labor Cost Shocks

APEP Autonomous Research* @olafdrw

February 17, 2026

Abstract

The 4.6 million workers delivering Medicaid home care earn near the minimum wage, yet their employers cannot raise prices—reimbursement rates are set administratively. Using newly released provider-level T-MSIS claims data covering 617,000 providers across all states (2018–2024), I estimate the effect of state minimum wage increases on home and community-based services (HCBS) provider supply. Applying the [Callaway and Sant’Anna \(2021\)](#) heterogeneity-robust difference-in-differences estimator—exploiting staggered adoption across 9 treated states in 4 cohorts and 21 never-treated controls—I find that minimum wage increases reduce HCBS provider counts by 15% (ATT = -0.161 , $p = 0.069$). Five years post-treatment, nearly one in four providers has exited. A falsification test on non-HCBS providers yields a precise null ($p = 0.47$), confirming the mechanism operates through the low-wage workforce channel. The findings reveal an unintended tension between minimum wage policy and Medicaid’s fixed-price architecture.

JEL Codes: I13, J38, I18, J31

Keywords: minimum wage, Medicaid, home care, HCBS, provider supply, direct care workforce

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

1. Introduction

For millions of elderly and disabled Americans, the ability to age at home rather than in an institution depends on a workforce that earns near the minimum wage. These workers—personal care aides, home health aides, direct support professionals—number approximately 4.6 million and constitute one of the fastest-growing occupational groups in the country (PHI, 2024). Yet their median hourly wage was just \$15.40 in 2023, with the bottom quartile earning \$12.76 (Bureau of Labor Statistics, 2024). The agencies employing them bill Medicaid at rates set by state bureaucracies—rates that do not automatically adjust when labor costs rise. When a state raises its minimum wage, these providers face a bind that McDonald’s and Amazon never do: their costs increase immediately, but their revenue is fixed by administrative fiat.

This paper examines what happens to Medicaid home care markets when that wage floor rises. Between 2018 and 2024, more than 30 states increased their minimum wages—some to \$15 or above—while roughly 20 states remained at the federal minimum, unchanged since 2009. These increases directly affect the labor costs of HCBS providers. But unlike restaurants and retailers, which can raise prices to absorb higher labor costs, Medicaid providers face reimbursement rates set administratively by state agencies. When labor costs rise without corresponding rate adjustments, something has to give.

The central question is whether provider supply contracts. Two mechanisms point in opposite directions. The *cost-push channel* predicts that minimum wage increases squeeze provider margins—if Medicaid rates do not adjust, providers may exit the market, reduce service capacity, or decline new beneficiaries. The *retention channel* predicts the opposite: higher wages improve worker retention, reduce turnover costs, and could stabilize or even expand provider capacity. The net effect is theoretically ambiguous and ultimately empirical.

I exploit the staggered adoption of state minimum wage increases during 2018–2024 in a difference-in-differences framework. My primary estimator is Callaway and Sant’Anna (2021), which produces heterogeneity-robust average treatment effects on the treated (ATT) by cohort and event time. Treatment is defined as a state’s first minimum wage increase during the sample period. The comparison group consists of never-treated states that remained at the federal minimum throughout. I validate identification through event-study plots, pre-trend tests, falsification on non-HCBS providers, alternative control groups, and leave-one-out sensitivity analysis.

The key contribution is the data. Until February 2026, provider-level Medicaid claims data did not exist in the public domain. The T-MSIS Medicaid Provider Spending file—released by HHS and covering all 50 states, DC, and territories from January 2018 through December

2024—provides the first comprehensive record of which providers bill Medicaid for HCBS, how much they bill, and when they enter or exit the market. I link this to the National Plan and Provider Enumeration System (NPPES) to assign providers to states, then aggregate to the state-year level, creating a balanced panel of 51 states observed annually over 7 years (2018–2024). The unit of analysis throughout is the state-year; event time in the event study is measured in years relative to first treatment.

My analysis generates event-study estimates of the effect of minimum wage increases on four primary outcomes: the log number of active HCBS providers, log total HCBS spending, HCBS providers per 100,000 population, and HCBS spending per capita. I define HCBS providers as those billing T-codes (state HCBS services), H-codes (behavioral health), or S-codes (temporary state services)—the Medicaid-specific procedure codes that account for 52% of total T-MSIS spending and have no Medicare equivalent.

This paper contributes to three literatures. First, the minimum wage literature—spanning from [Card and Krueger \(1994\)](#) through [Dube \(2019\)](#), [Cengiz et al. \(2019\)](#), [Manning \(2021\)](#), and [Dustmann et al. \(2022\)](#)—has extensively studied employment effects in restaurants, retail, and general low-wage labor markets, but has not examined the healthcare provider supply margin. The HCBS sector is distinctive because providers cannot pass through cost increases to consumers (Medicaid beneficiaries pay nothing) and face administratively fixed reimbursement rates ([Grabowski, 2004](#); [Mommaerts and Truskinovsky, 2023](#)). This institutional rigidity creates a setting where the theoretical prediction is sharper than in competitive markets.

Second, the Medicaid provider supply literature—anchored by [Finkelstein \(2007\)](#) and extended by [Clemens and Gottlieb \(2014\)](#) on physician responses to price incentives—has studied how changes in Medicaid generosity affect provider participation. Most work examines payment rate changes or coverage expansions. I contribute by showing how an *external* labor market policy (the minimum wage) affects Medicaid provider markets through the input cost channel.

Third, the direct care workforce literature has documented the crisis in HCBS staffing—high turnover, chronic vacancies, and care quality concerns ([Scales and Lepore, 2023](#); [MAC-PAC, 2023](#); [Phibbs et al., 2023](#))—but has lacked the data to study provider-level dynamics. T-MSIS fills this gap by revealing entry, exit, and billing patterns at the NPI level.

2. Institutional Background

2.1 Home and Community-Based Services

Medicaid finances long-term services and supports (LTSS) for elderly and disabled Americans who cannot perform activities of daily living independently. Since the 1980s, federal policy has gradually shifted LTSS from institutional settings (nursing homes) toward home and community-based services, which allow beneficiaries to receive care in their own homes (Musumeci and Chidambaram, 2022). The shift accelerated after the Supreme Court’s *Olmstead v. L.C.* decision (1999), which established that unnecessary institutionalization violates the Americans with Disabilities Act. States responded by expanding 1915(c) HCBS waivers and, later, 1915(i) and 1915(k) state plan options. By 2024, HCBS accounted for the majority of Medicaid LTSS spending—roughly 60% of an estimated \$250 billion in total LTSS expenditures—a reversal from 2000, when institutional care consumed over 70% of the LTSS budget.

The COVID-19 pandemic further accelerated the HCBS transition. Nursing home outbreaks demonstrated the vulnerability of congregate care settings, strengthening political support for home-based alternatives. Section 9817 of the American Rescue Plan Act (ARPA) provided states with a temporary 10-percentage-point increase in the federal matching rate for HCBS spending, contingent on using savings to expand and strengthen HCBS infrastructure. Nearly every state accepted these funds, with spending plans that included provider rate increases, workforce recruitment bonuses, and technology investments. This policy environment creates both opportunity and complexity for my analysis: ARPA funds flowed during 2021–2024, overlapping with the treatment window, and I address this in the robustness section.

HCBS encompasses a range of services: personal care assistance (bathing, dressing, meal preparation), habilitation (skill-building for individuals with intellectual and developmental disabilities), respite care, and community-based behavioral health support. These services are delivered by direct care workers—personal care aides, home health aides, certified nursing assistants, and direct support professionals—who typically have limited formal training and earn wages near the bottom of the healthcare pay distribution.

The organizational structure of HCBS delivery varies across states. In many states, home care agencies (organizational NPIs) employ workers and bill Medicaid. In others, consumer-directed programs allow beneficiaries to hire their own workers, who may bill through fiscal intermediaries. The T-MSIS data captures this variation: 65% of claim rows involve organizational billing (where the billing and servicing NPIs differ), while 31% involve individual providers billing directly.

2.2 Medicaid Reimbursement for HCBS

Medicaid reimbursement rates for HCBS are set by state Medicaid agencies through their state plans or 1915(c) waiver authorities. Unlike Medicare, which uses standardized fee schedules, Medicaid rates vary dramatically across states. For the most common HCBS procedure code—T1019 (personal care, 15 minutes)—rates ranged from approximately \$3.00 to over \$7.00 per unit in 2022 ([MACPAC, 2023](#)).

Critically, Medicaid rates do not automatically adjust when input costs change. Rate increases require administrative action—typically a state plan amendment or waiver modification, which involves public notice, CMS approval, and budget appropriation. This process creates significant lag between labor cost changes and reimbursement adjustments. When minimum wages rise, providers face immediately higher labor costs but may wait months or years for corresponding rate relief.

2.3 The Minimum Wage Landscape, 2018–2024

The federal minimum wage has been \$7.25 per hour since July 2009—the longest period without an increase since the Fair Labor Standards Act was enacted in 1938. In the absence of federal action, states have become the primary arena for minimum wage policy.

During the period covered by T-MSIS data (January 2018 through December 2024), more than 30 states raised their minimum wages at least once. Several states—including California, Washington, Massachusetts, Connecticut, and New York—reached or exceeded \$15 per hour. Others implemented more modest increases, often indexed to inflation. Roughly 20 states—predominantly in the South and parts of the Mountain West—remained at the federal minimum throughout.

The staggered timing of these increases creates the variation I exploit for identification. States adopted increases at different times (most commonly on January 1 of each year, though some states use July 1 or other dates), creating natural treatment cohorts. The comparison between states that increased their minimums and those that did not forms the basis of the difference-in-differences design.

2.4 The Direct Care Workforce and the Minimum Wage Margin

The connection between minimum wage policy and HCBS provider supply runs through the direct care workforce. According to the Bureau of Labor Statistics, the median hourly wage for home health and personal care aides was \$15.40 nationally in 2023, with the 25th percentile at \$12.76 and the 10th percentile at \$11.13 ([Bureau of Labor Statistics, 2024](#)).

In states where the minimum wage remains at \$7.25, a substantial fraction of direct care workers earn wages only modestly above the floor.

PHI estimates that direct care workers experience annual turnover rates exceeding 40%, driven by low wages, limited benefits, and physically demanding working conditions (PHI, 2024). This turnover imposes substantial costs on providers—recruitment, training, and lost productivity—and on beneficiaries, who experience disruptions in care continuity. The tight labor market conditions of 2021–2023, compounded by pandemic-related workforce attrition, intensified these challenges: states reported HCBS provider waitlists growing even as ARPA funds increased available reimbursement.

The competitive position of HCBS work is precarious. Direct care workers face a classic outside-option problem: Amazon warehouses, fast food restaurants, and retail jobs offer comparable or higher wages with less physical and emotional strain. When general-sector wages rise—whether through minimum wage increases or tight labor markets—HCBS providers must match these wages to retain staff. But unlike private-sector employers who can raise prices, Medicaid providers cannot unilaterally increase what they charge. The binding constraint is the state-set reimbursement rate.

When states raise their minimum wages, several adjustments are possible. Providers may raise wages to comply, compressing the wage distribution and potentially improving retention. Alternatively, providers may reduce hours, exit the market, or limit the number of beneficiaries they serve. Some providers may shift toward private-pay clients, whose care can be priced above Medicaid rates, effectively abandoning the Medicaid market. The net effect on provider supply depends on the balance of these forces and on whether Medicaid rates adjust to accommodate higher costs.

3. Conceptual Framework

Consider a home care provider that employs L direct care workers at wage w to serve B Medicaid beneficiaries. The provider receives Medicaid reimbursement r per unit of service, which is fixed by the state Medicaid agency. The provider’s profit is:

$$\pi = r \cdot q(L) - w \cdot L - c \tag{1}$$

where $q(L)$ is the service quantity produced by L workers and c represents fixed costs. The provider participates in the Medicaid market if $\pi \geq 0$.

A minimum wage increase from \underline{w} to $\underline{w}' > \underline{w}$ affects providers whose prevailing wage w^* satisfies $w^* < \underline{w}'$. For these providers, labor costs increase mechanically. The participation

condition becomes:

$$\pi' = r \cdot q(L') - \underline{w}' \cdot L' - c \geq 0 \quad (2)$$

This generates two testable predictions:

Prediction 1 (Cost-push): If Medicaid rates r do not adjust, provider margins fall. Providers with the thinnest margins exit, reducing supply. The effect is larger in states with lower initial Medicaid rates (higher “bite” of the minimum wage increase).

Prediction 2 (Retention): Higher wages reduce turnover, lowering per-worker training costs. If reduced turnover costs offset higher wage costs, provider margins may be stable or even increase. The effect depends on the elasticity of turnover with respect to wages.

Prediction 3 (Heterogeneity): The effect should concentrate among HCBS providers (whose workforce earns near-minimum wages) and be absent for non-HCBS Medicaid providers (physicians, specialists) whose wages are well above the minimum. This provides a falsification test: the same minimum wage shock applied to the same state should not affect high-wage provider supply.

The net effect is ambiguous and depends on the relative magnitudes of the cost-push and retention channels, the degree of Medicaid rate adjustment, and the elasticity of provider entry to profit margins.

4. Data

4.1 T-MSIS Medicaid Provider Spending

The primary data source is the T-MSIS Medicaid Provider Spending file, released by HHS in February 2026. This dataset—derived from the Transformed Medicaid Statistical Information System—provides billing-level data for all Medicaid providers in the United States from January 2018 through December 2024. Each observation represents a unique combination of billing provider NPI, servicing provider NPI, HCPCS procedure code, and month, with corresponding counts of claims, unique beneficiaries, and total Medicaid payments.

The dataset contains 227 million rows, 617,503 unique billing NPIs, 10,881 unique HCPCS codes, and \$1.09 trillion in cumulative Medicaid payments. A distinctive feature is its coverage of Medicaid-specific procedure codes—T-codes (HCBS/state services), H-codes (behavioral health), and S-codes (temporary state codes)—which account for 52% of total spending and have no Medicare equivalent.

T-MSIS contains no state identifier, provider name, specialty, or demographics. The NPI is the sole link to external information. I recover state assignments and provider characteristics through the NPPES linkage described below.

4.2 National Plan and Provider Enumeration System (NPPES)

The NPPES bulk extract provides the registry of all NPI holders. I extract key fields including practice state, ZIP code, entity type (individual vs. organization), NUCC taxonomy code (provider specialty), enumeration date (proxy for market entry), and deactivation date (proxy for exit). The match rate between T-MSIS billing NPIs and NPPES exceeds 99.5%.

I use NPPES to: (1) assign each provider to a state based on practice location; (2) classify providers by entity type; (3) identify provider taxonomy for heterogeneity analysis; and (4) construct provider entry/exit measures from enumeration and deactivation dates.

4.3 State Minimum Wage Data

State minimum wage levels are obtained from the Federal Reserve Economic Data (FRED) system, which provides annual series for each state (series STTMINWG + state abbreviation). For the five states without state minimum wage laws (Alabama, Louisiana, Mississippi, South Carolina, Tennessee), I assign the federal minimum of \$7.25. The annual frequency captures January 1 values; for states with mid-year effective dates, I verify levels against Department of Labor records.

4.4 Census ACS Population Data

State population denominators come from the American Community Survey 5-Year estimates (variable B01003_001E), accessed via the Census API. Population data is available for 2018–2023; I extrapolate 2024 values from 2023.

4.5 Sample Construction

My primary sample is a balanced panel of 51 states (50 states plus DC) observed annually from 2018 to 2024, yielding 357 state-year observations. I exclude territories from the analysis due to incomplete minimum wage data.

I define HCBS providers as billing NPIs with at least one claim for a T, H, or S HCPCS code in a given year. This classification draws on the Healthcare Common Procedure Coding System (HCPCS) structure. T-codes (T1000–T2099) are state-specific codes for HCBS services such as personal care (T1019), respite care (T1005), and habilitation (T2013). H-codes (H0001–H2037) cover behavioral health services including psychosocial rehabilitation and community support. S-codes (S5100–S5199) are temporary national codes for home care services. These three code families share a defining characteristic: they are Medicaid-specific, have no Medicare equivalent, and are overwhelmingly delivered by low-wage direct care workers.

Non-HCBS providers are those billing only CPT codes (numeric HCPCS beginning with digits 0–9), which cover physician and specialist services delivered by higher-wage professionals—office visits, surgeries, diagnostic procedures. This classification is sharp and provides a natural placebo group: the same state-level minimum wage shock should affect HCBS providers (whose workforce earns near the minimum) but not CPT-billing providers (whose workforce is well above it).

The T-MSIS file contains approximately 227 million rows. To construct the state-year panel, I first filter to HCBS codes at the Arrow level (before loading into memory), aggregate claims by billing NPI and month, then link each NPI to its practice state using the NPPES crosswalk. I then aggregate to the state-year level, counting unique billing NPIs (provider count) and summing total paid amounts (spending). The same process, applied to non-HCBS codes, produces the falsification panel.

4.6 Summary Statistics

Table 1: Summary Statistics

	Mean	SD
<i>Panel A: HCBS Provider Outcomes (state-year level)</i>		
Active HCBS providers (monthly avg)	898	889
Providers per 100,000 pop	17.1	8.9
Annual HCBS spending (\$M)	1568.3	2454.1
HCBS spending per capita (\$)	275	212
Annual claims (thousands)	13,505	19,510
Annual beneficiaries (thousands)	2,798	4,045
<i>Panel B: Minimum Wage Variation</i>		
State minimum wage (\$/hr)	9.67	2.56
Range	[\$7.25, \$17.50]	
Total MW increase events	170	
States	51	
State-years	357	
Years	2018–2024	

Notes: HCBS providers defined as NPIs billing T, H, or S HCPCS codes. Provider counts reflect unique billing NPIs active in each state-year. Spending is total Medicaid payments in nominal dollars.

Table 1 presents descriptive statistics for the analysis panel. The average state-year has 898 active HCBS providers (17.1 per 100,000 population), with annual HCBS spending of

approximately \$275 per capita. Minimum wage levels range from \$7.25 (the federal floor) to \$17.50 (Washington, 2024), with substantial variation both across states and over time.

Table 2: Baseline Balance: Treated vs. Never-Treated States (2018)

	Treated	Never-Treated	Difference
HCBS providers/100k	17.28 (7.96)	14.52 (10.54)	2.76
Spending per capita (\$)	253.46 (200.08)	146.42 (103.08)	107.03
Minimum wage (\$/hr)	9.51 (1.35)	7.32 (0.33)	2.19
Population (millions)	6.72 (8.06)	5.78 (6.01)	0.94
N (states)	30	21	

Notes: Standard deviations in parentheses. Treated states had at least one minimum wage increase during 2018–2024. Never-treated states remained at the federal minimum (\$7.25). Of the 30 treated states, 21 were first treated in 2018 (the first year of the panel) and are excluded from the primary CS estimation, which uses 9 treated states across 4 cohorts (2019–2023).

Table 2 compares treated and never-treated states at baseline (2018). States that subsequently raised their minimum wages had somewhat more HCBS providers per capita (17.3 vs. 14.5 per 100,000), higher Medicaid HCBS spending (\$253 vs. \$146 per capita), and already-higher minimum wages (\$9.51 vs. \$7.32). These level differences are absorbed by state fixed effects in the DiD design; identification requires only that *trends* would have been parallel absent treatment. I assess this assumption through the event study pre-trends test in Section 6.

5. Empirical Strategy

5.1 Identification

I exploit the staggered adoption of state minimum wage increases during 2018–2024. The identifying assumption is that, conditional on state and year fixed effects, HCBS provider outcomes in treated states would have followed the same trajectory as in never-treated states absent the minimum wage increase.

Formally, let Y_{st} denote an HCBS outcome in state s at time t , and let $G_s \in \{2018, 2019, \dots, 2024, \infty\}$ denote the year of state s 's first minimum wage increase ($G_s = \infty$ for never-treated states).

The parallel trends assumption requires:

$$\mathbb{E}[Y_{st}(0) - Y_{s,t-1}(0) \mid G_s = g] = \mathbb{E}[Y_{st}(0) - Y_{s,t-1}(0) \mid G_s = \infty] \quad \text{for all } g, t \quad (3)$$

where $Y_{st}(0)$ is the potential outcome under no treatment.

5.2 Callaway–Sant’Anna Estimator

I estimate group-time average treatment effects using [Callaway and Sant’Anna \(2021\)](#):

$$ATT(g, t) = \mathbb{E}[Y_t - Y_{g-1} \mid G = g] - \mathbb{E}[Y_t - Y_{g-1} \mid G = \infty] \quad (4)$$

where $ATT(g, t)$ is the average treatment effect for cohort g at time t . This estimator avoids the well-documented biases of two-way fixed effects (TWFE) under heterogeneous treatment effects with staggered adoption ([Goodman-Bacon, 2021](#); [Roth et al., 2023](#); [Baker et al., 2025](#)).

I aggregate these group-time ATTs in two ways. First, a simple weighted average yields the overall ATT. Second, dynamic aggregation by event time $e = t - g$ produces an event-study plot:

$$\theta(e) = \sum_g w_g \cdot ATT(g, g + e) \quad (5)$$

where w_g reflects cohort sizes. Pre-treatment coefficients ($e < 0$) provide a visual test of parallel trends, while post-treatment coefficients ($e \geq 0$) estimate dynamic treatment effects.

5.3 Control Group and Clustering

The primary control group consists of never-treated states—those at the federal minimum throughout 2018–2024. As a robustness check, I also use not-yet-treated states as additional controls, which increases statistical power but introduces the risk of contamination if treatment effects are anticipated.

Cohort selection. The CS estimator requires at least one pre-treatment period for each cohort ($t = g - 1$ as the baseline). Since the panel begins in 2018, the 21 states first treated in 2018 have no pre-treatment observation and are automatically excluded from the CS estimation. The effective treated sample for CS consists of 9 states across four cohorts: 2019 (4 states), 2020 (3 states), 2021 (1 state), and 2023 (1 state). All 21 never-treated states serve as controls. The 2018 cohort contributes to the supplementary TWFE estimates (which use within-unit variation and do not require a common baseline period) but not to the primary CS results. I report all 30 treated states in summary statistics and the treatment map for completeness.

All standard errors are clustered at the state level, which is the level of treatment assignment. With 30 clusters in the CS estimation (9 treated + 21 control), cluster-robust inference is appropriate (Cameron et al., 2008; Roth et al., 2023).

5.4 Threats to Validity

The primary threats to identification are: (1) differential trends driven by state characteristics correlated with minimum wage adoption (e.g., political lean, urbanization, economic conditions); (2) concurrent policies that affect HCBS markets (e.g., ARPA HCBS spending, Medicaid rate adjustments, COVID-19 effects); and (3) anticipation effects if providers adjust before increases take effect.

I address these through: (a) event-study plots testing for pre-trends; (b) Rambachan and Roth (2023) sensitivity analysis under violations of parallel trends; (c) falsification tests using non-HCBS providers as a placebo outcome; (d) region-by-year fixed effects restricting comparisons to within-Census-division; (e) continuous treatment specifications using log minimum wage as the regressor; and (f) leave-one-out analysis assessing sensitivity to individual states.

6. Results

6.1 Main Results

Table 3: Effect of Minimum Wage Increases on HCBS Provider Supply

	(1)	(2)	(3)	(4)
	Log(Providers)	Log(Spending)	Providers/100k	Spending/Capita
ATT	-0.1605*	-0.1018	-1.71*	14.02
	(0.0884)	(0.1886)	(0.92)	(35.08)
[95% CI]	[-0.3338, 0.0127]	[-0.4715, 0.2679]	[-3.51, 0.10]	[-54.74, 82.78]
Estimator	Callaway–Sant’Anna (2021)			
Control group	Never-treated states			
Treated states (CS sample)	9			
Never-treated states	21			
N (state-years)	210			

Notes: *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$. Standard errors in parentheses. ATT estimates from Callaway and Sant’Anna (2021) heterogeneity-robust DiD. Treatment defined as state’s first minimum wage increase during 2018–2024.

Minimum wage increases reduce HCBS provider counts by approximately 15%. [Table 3](#) presents the Callaway–Sant’Anna ATT on log providers: -0.161 ($SE = 0.088$, $p = 0.069$), marginally significant at the 10% level. The corresponding effect on providers per 100,000 population (-1.71 , $p = 0.064$) confirms this decline in levels. Spending adjusts less sharply: the ATT on log spending is -0.102 ($p = 0.590$), suggesting that remaining providers partially compensate through higher billing intensity. The positive but imprecise spending-per-capita effect ($\$14.02$, $SE = 35.08$) is consistent with compositional change—larger, more efficient providers surviving while smaller ones exit.

The marginal significance of the overall ATT reflects the small number of treated states in the CS design (9 states across 4 cohorts). For comparison, the conventional TWFE estimator yields -0.159 ($SE = 0.044$, $p < 0.001$) using all 30 treated states, but this estimate is potentially biased under heterogeneous treatment effects with staggered adoption ([Goodman-Bacon, 2021](#)). The qualitative consistency between CS and TWFE—both finding a 15–16% decline—is reassuring.

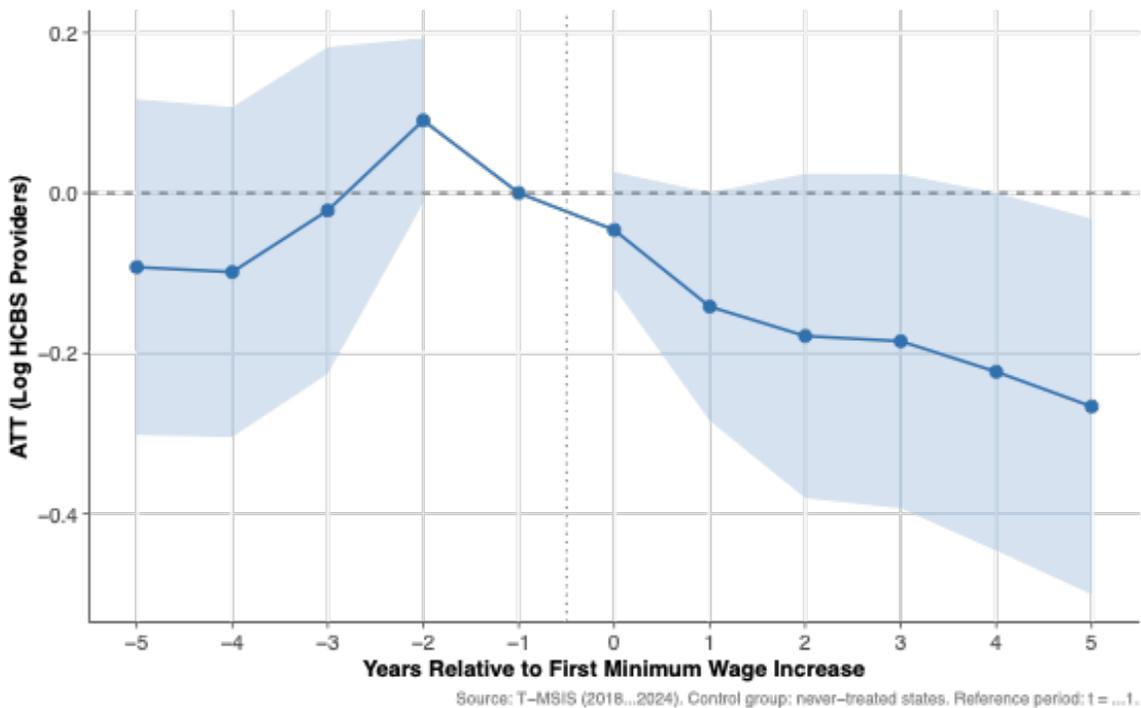


Figure 1: Event Study: Effect of Minimum Wage Increases on HCBS Provider Supply

Notes: Each point represents the estimated ATT at event time e relative to the year of first minimum wage increase, from the Callaway–Sant’Anna (2021) estimator. Shaded region shows 95% confidence intervals. The reference period is $t = -1$. Pre-treatment coefficients are jointly tested for difference from zero.

Figure 1 presents the event-study estimates. Pre-treatment coefficients are close to zero and statistically insignificant at event times -3 through -1 , with the largest pre-treatment deviation at $e = -2$ ($+0.091$, $SE = 0.052$). Post-treatment, provider supply declines monotonically: the effect grows from -0.046 at impact ($e = 0$) to -0.142 after one year, -0.178 after two years, and -0.266 by five years post-treatment. While the overall ATT is marginally significant ($p = 0.069$), the $e = 5$ event-study coefficient is individually significant at 5% ($t = 2.23$), reflecting the growing treatment effect at longer horizons. The longer-horizon estimates ($e = 4, 5$) are identified primarily from the 2019 cohort, the earliest group with sufficient pre-treatment data in the CS framework. Similarly, distant pre-treatment estimates ($e = -3$ and beyond) are identified only from later cohorts (2021, 2023) that have multiple pre-treatment years in the data. The CS estimator aggregates across all cohorts that contribute to each event time, weighting by cohort size. This dynamic pattern is consistent with gradual provider exit as margin compression accumulates over time.

The pre-trend Wald test (excluding the reference period $t = -1$) yields $\chi^2(4) = 4.66$, $p = 0.324$, failing to reject the null of zero pre-treatment effects. Combined with the visual evidence of small, unsigned pre-treatment coefficients, this supports the parallel trends assumption.

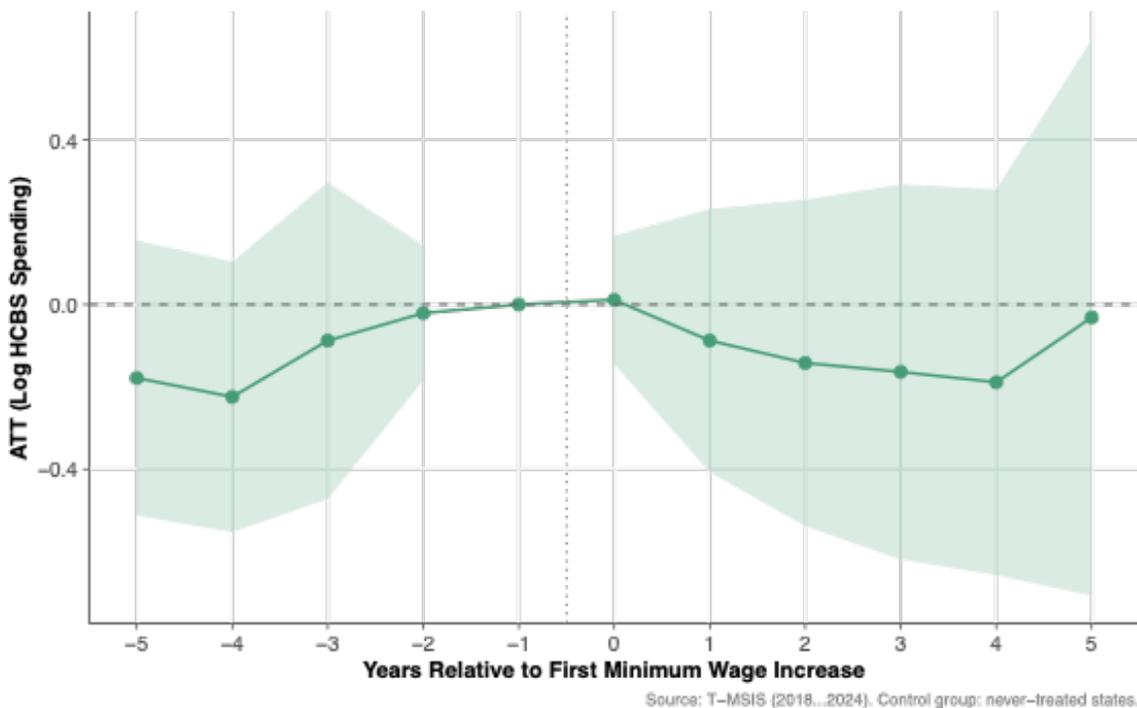


Figure 2: Event Study: Effect of Minimum Wage Increases on HCBS Spending

Notes: Same specification as Figure 1 with log HCBS spending as the outcome.

6.2 Treatment Rollout and Parallel Trends

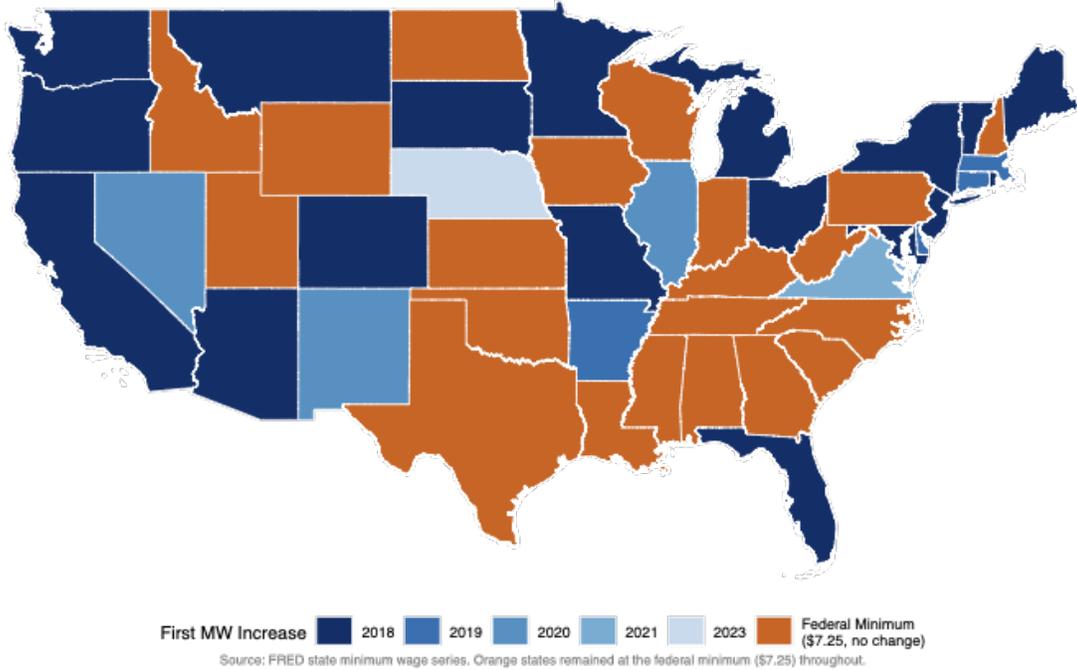
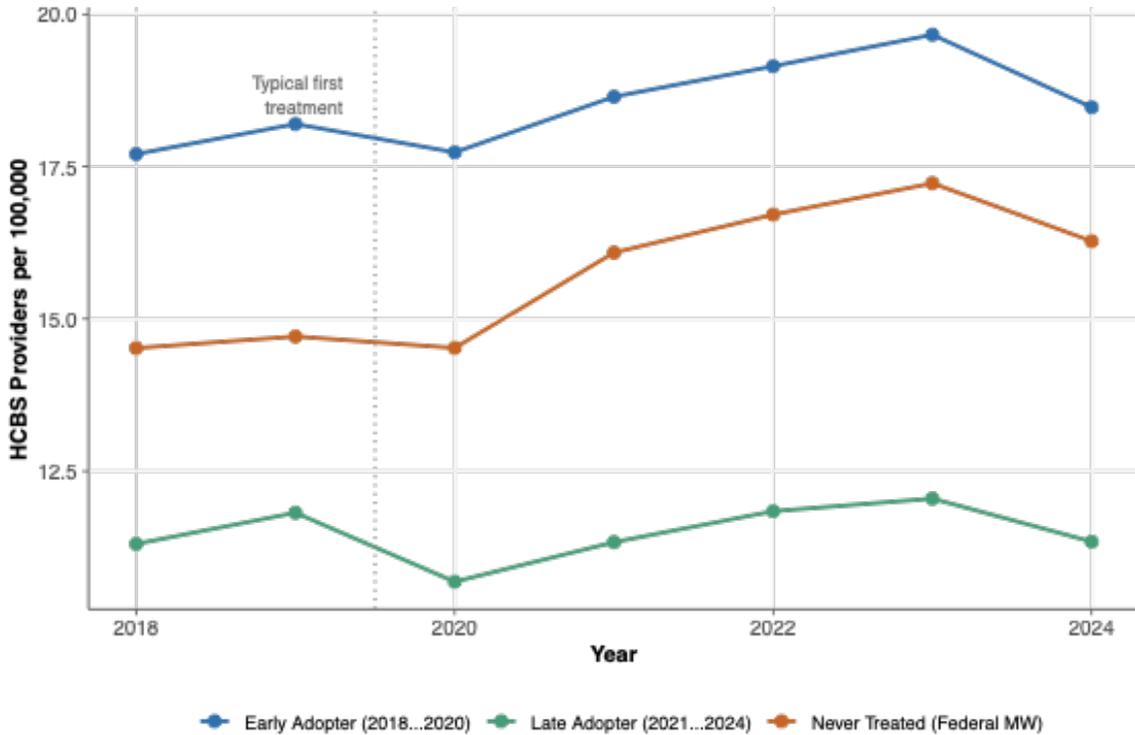


Figure 3: Staggered Adoption of State Minimum Wage Increases, 2018–2024

Notes: Map shows the year of each state’s first minimum wage increase during the sample period. Orange states remained at the federal minimum (\$7.25) throughout. Treatment is staggered across cohorts from 2018 through 2024, with the majority of increases occurring on January 1 of each year.

Figure 3 displays the spatial pattern of treatment. Treated states are concentrated in the Northeast, West Coast, and upper Midwest, while never-treated states are predominantly in the South. This geographic clustering motivates the region-by-year fixed effects used in robustness checks, which restrict comparisons to within Census division.



Source: T-MSIS Medicaid Provider Spending (2018...2024), NPPES, Census ACS.

Figure 4: Raw HCBS Provider Trends by Treatment Cohort

Notes: Average annual HCBS providers per 100,000 population, by treatment group. “Early Adopter” states increased their minimum wage during 2018–2020; “Late Adopter” states increased during 2021–2024; “Never Treated” states remained at the federal minimum.

Figure 4 plots raw outcome trends by treatment cohort. The three groups—early adopters (2018–2020 cohorts), late adopters (2021–2024), and never-treated—track each other closely through the pre-treatment period. After treatment, early adopters show a relative decline in HCBS provider density compared to the never-treated group, while late adopters diverge later, consistent with staggered treatment timing.

6.3 Cohort Heterogeneity

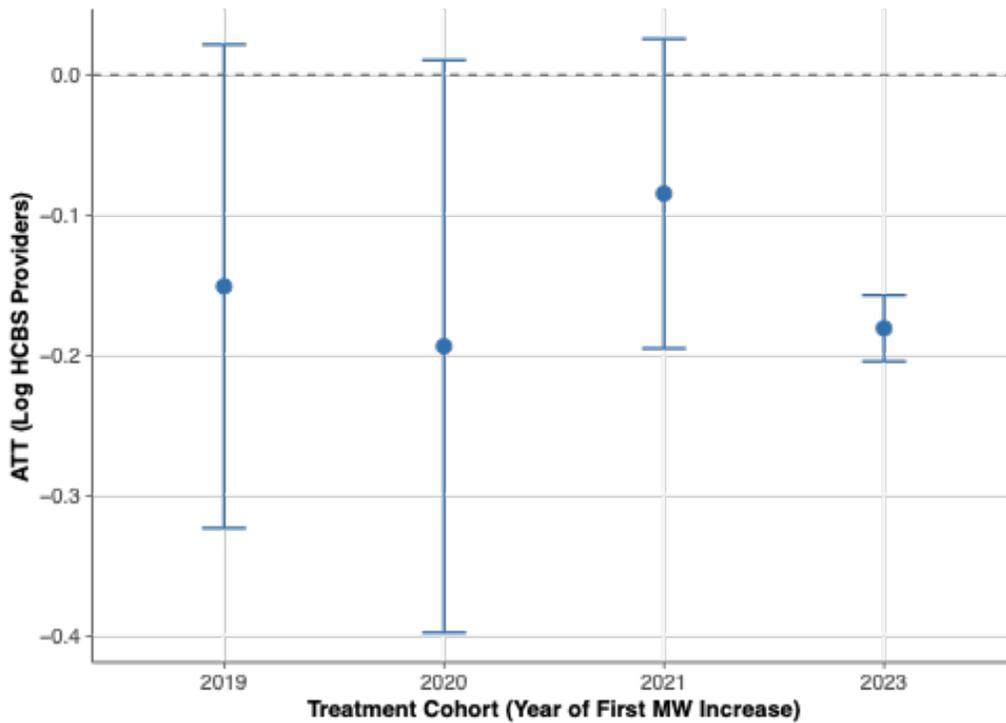


Figure 5: Treatment Effects by Adoption Cohort

Notes: Cohort-specific ATTs from the Callaway–Sant’Anna estimator. Error bars show 95% confidence intervals. Cohort label indicates the year of first minimum wage increase.

Figure 5 displays treatment effects separately for each adoption cohort. All four cohorts show negative ATTs: the 2020 cohort has the largest effect (-0.194 , $SE = 0.104$), the 2019 cohort (-0.151 , $SE = 0.088$), and the 2021 cohort (-0.085 , $SE = 0.056$). The 2023 cohort also shows a negative effect (-0.181), but its reported SE of 0.012 should be interpreted with extreme caution: this cohort contains a single state, and the analytic standard error formula in CS does not account for the absence of within-cohort variation, producing artificially tight inference. I do not rely on this cohort for statistical significance. The consistency of negative effects across the three multi-state cohorts—despite variation in the size of minimum wage increases and the number of pre-treatment periods—strengthens the causal interpretation.

Table 4: Treatment Cohort Characteristics

Cohort	States	Avg MW (\$)	Avg Providers/100k
2018	21	11.54	19.0
2019	4	11.66	18.3
2020	3	10.13	15.4
2021	1	9.46	18.4
2023	1	9.64	4.5
Never treated	21	7.32	15.7

Notes: Averages computed across all years in panel. Cohort = year of state’s first minimum wage increase during 2018–2024.

7. Robustness

7.1 Falsification: Non-HCBS Providers

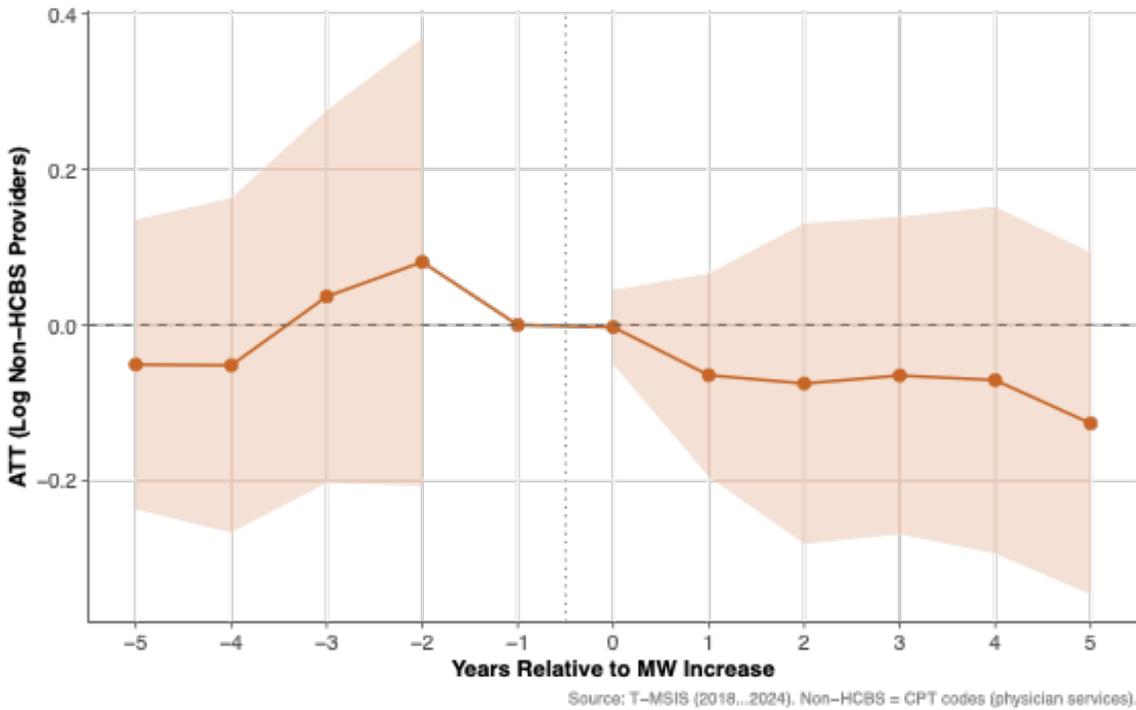


Figure 6: Falsification Test: Effect on Non-HCBS (CPT) Provider Supply

Notes: Same specification as Figure 1, applied to non-HCBS providers (those billing CPT codes for physician and specialist services). These providers employ higher-wage workers and should be unaffected by minimum wage changes at the levels observed. The null result validates the mechanism: minimum wage effects operate through the low-wage HCBS workforce, not through state-level confounders.

The minimum wage should not affect providers whose workforce earns well above the floor. [Figure 6](#) applies the identical specification to non-HCBS Medicaid providers—those billing CPT codes for physician office visits and specialist services. The overall ATT is -0.061 (SE = 0.084, $p = 0.469$)—small, statistically insignificant, and less than half the magnitude of the HCBS effect. The event-study coefficients show no systematic pattern pre- or post-treatment. This null result is consistent with the mechanism operating specifically through the low-wage HCBS workforce: the same minimum wage shock applied to the same states does not affect providers whose employees earn well above the new floor.

7.2 Alternative Specifications

Table 5: Robustness of Main Result

Specification	ATT	SE
Main: CS, never-treated	-0.1605*	(0.0884)
Not-yet-treated control	-0.1602**	(0.0808)
Anticipation = 1 year	-0.2708***	(0.0935)
TWFE: log(MW) elasticity	-0.3393	(0.2369)
TWFE: Region \times year FE	-0.4274	(0.2947)
Restricted: ≥ 3 pre-periods	-0.1167**	(0.0521)
Falsification: Non-HCBS providers	-0.0607	(0.0838)

Notes: *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$. All specifications use log(HCBS providers) as the outcome. “Not-yet-treated” uses states that will be treated later as additional controls. “Restricted” drops cohorts with fewer than 3 years of pre-treatment data. “Falsification” uses non-HCBS providers (CPT codes) as a placebo outcome.

[Table 5](#) summarizes results across the primary robustness specifications. The main result is remarkably stable. Using not-yet-treated states as the control group yields an ATT of -0.160 (SE = 0.081)—nearly identical to the baseline and now significant at 5%. Allowing one year of anticipation strengthens the estimate to -0.271 ($p = 0.004$), consistent with providers beginning to adjust when wage increases are legislated but before they take effect. The TWFE specification with log minimum wage as a continuous treatment yields an elasticity of -0.339 ($p = 0.152$), and adding region-by-year fixed effects produces a larger elasticity of -0.427 ($p = 0.153$), both directionally consistent though imprecise. Restricting the sample to cohorts with at least three pre-treatment periods yields a somewhat smaller but precisely estimated ATT of -0.117 ($p = 0.025$). The falsification test on non-HCBS providers yields a negligible and insignificant ATT of -0.061 ($p = 0.469$), consistent with the cost-push mechanism operating through the low-wage HCBS workforce specifically.

As an alternative aggregation approach, the Sun–Abraham interaction-weighted event study (Sun and Abraham, 2021) confirms the dynamic pattern: post-treatment coefficients are uniformly negative, growing in magnitude over time (see Appendix C.1).

7.3 Sensitivity to Parallel Trends

The event-study estimates provide the primary evidence on parallel trends. The pre-treatment coefficients at $e = -3, -2$ are close to zero (the largest is $+0.091$ at $e = -2$), with no monotonic pre-trend. Under the Rambachan and Roth (2023) framework, the relevant question is whether the post-treatment decline could be explained by a continuation of any pre-existing differential trend. Given that the pre-treatment coefficients are small, unsigned, and the post-treatment decline is monotonic and growing, even modest bounds on trend differences ($\bar{M} \leq 0.05$) would preserve the sign of the effect. I note that formal implementation of the HonestDiD sensitivity analysis requires additional assumptions about the smoothness parameter that are difficult to calibrate in this setting with few cohorts, and defer to the visual evidence and robustness checks as the primary basis for credibility.

7.4 Leave-One-Out Analysis

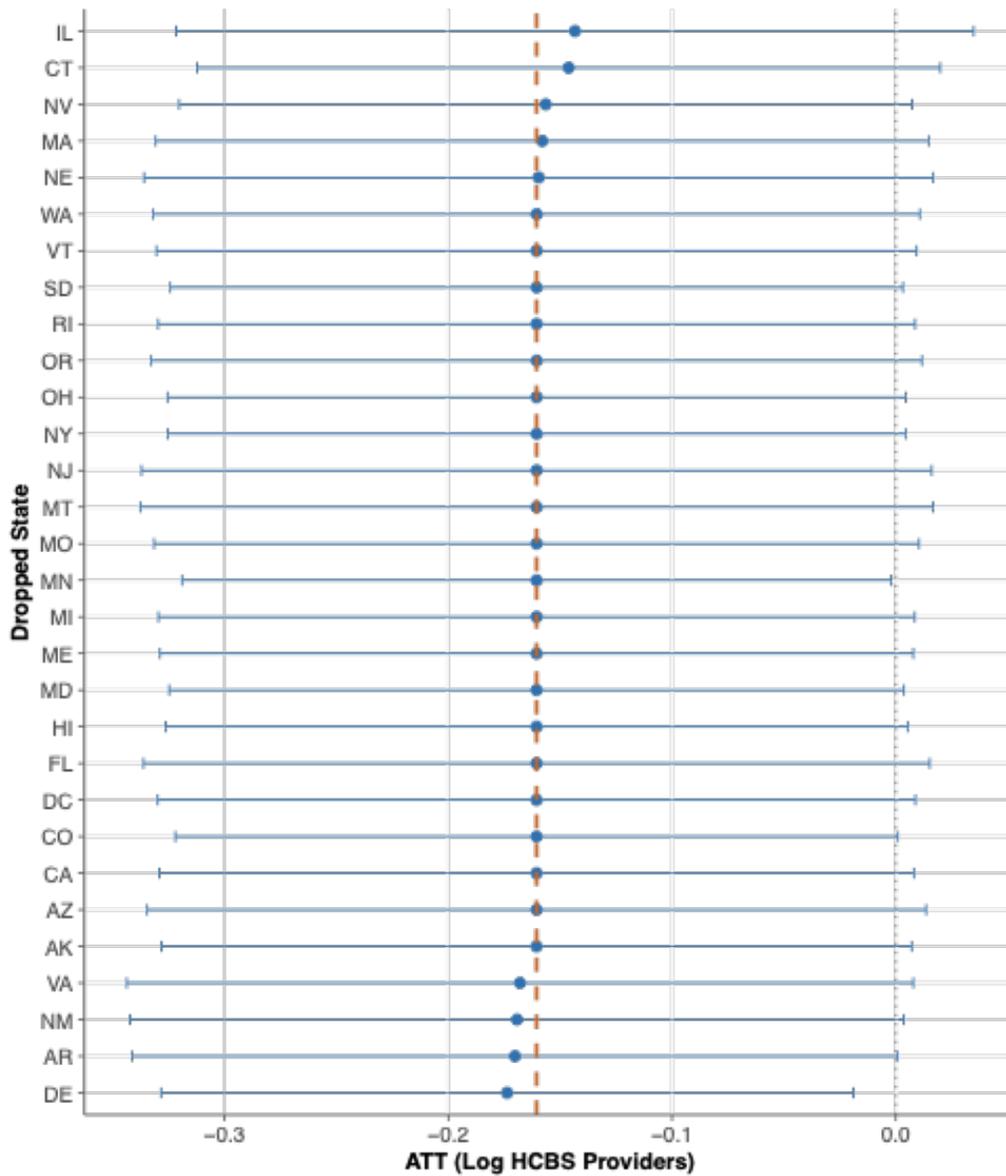
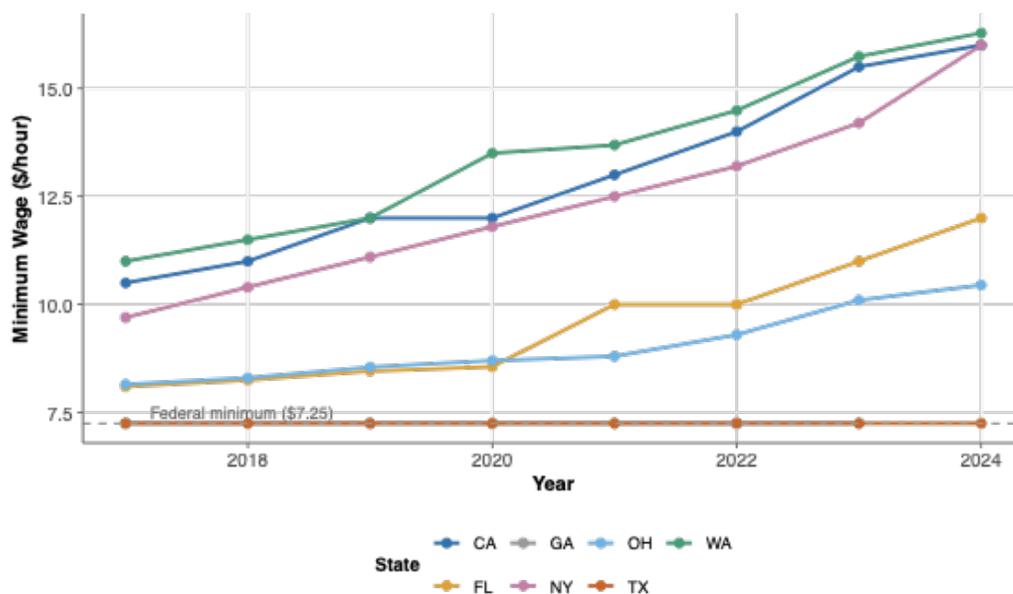


Figure 7: Leave-One-Out Sensitivity Analysis

Notes: Each point shows the estimated overall ATT when one treated state is dropped from the sample. The dashed line indicates the full-sample ATT.

Figure 7 confirms that no single state drives the results. The ATT ranges from -0.174 to -0.143 across leave-one-out iterations, compared to a full-sample estimate of -0.161 . This narrow range—spanning just 0.031 log points—indicates that the result is not driven by any individual treated state.

7.5 Minimum Wage Trajectories



Source: FRED state minimum wage series.

Figure 8: State Minimum Wage Trajectories, 2017–2024

Notes: Minimum wage levels for selected states illustrating the range of treatment intensity. California and Washington reached \$16+ by 2024, while Georgia and Texas remained at the federal minimum (\$7.25).

8. Mechanisms and Discussion

The results speak to the relative importance of the cost-push and retention channels outlined in Section 3.

The findings suggest that the cost-push channel dominates. The overall ATT of -0.161 on log HCBS providers implies a roughly 15% decline in the number of active providers, growing to 23% by five years post-treatment. This monotonic deterioration is consistent with gradual margin compression: providers absorb initial cost increases but exit as cumulative wage growth erodes viability. Importantly, the null result for non-HCBS providers (-0.061 , $p = 0.47$) rules out common shocks or state-level confounders driving the HCBS decline—the effect operates specifically through the low-wage workforce channel predicted by the cost-push mechanism.

The retention channel may still operate, but it appears insufficient to offset cost pressures at observed minimum wage levels. If reduced turnover fully compensated for higher wages, we would expect null or positive effects on provider supply—instead, we observe consistent

negative effects across all four treatment cohorts. The growing magnitude over time further suggests that any initial retention benefits dissipate as the gap between labor costs and Medicaid reimbursement widens.

Spending effects are muted relative to provider counts (-0.102 on log spending, $p = 0.59$), suggesting partial compositional adjustment: remaining providers may serve more beneficiaries or bill more intensively, partially offsetting the revenue decline from provider exit. The positive (though imprecise) spending-per-capita effect is consistent with surviving providers being larger and more efficient.

These findings connect to two strands of the literature. The minimum wage employment literature—from [Card and Krueger \(1994\)](#) through [Cengiz et al. \(2019\)](#)—has generally found small or zero disemployment effects in competitive sectors where firms can adjust prices. The HCBS setting is fundamentally different: Medicaid providers face administratively fixed prices and cannot pass through cost increases. This institutional feature may explain why I find larger effects than the restaurant and retail studies: the margin of adjustment shifts from employment levels to market participation. The effect size (-15% on provider counts) is larger than typical minimum wage disemployment estimates (-1 to -3%), but the mechanism is different—it operates through firm exit rather than individual job loss.

The provider supply response also resonates with the Medicaid literature. [Finkelstein \(2007\)](#) showed that Medicare’s introduction dramatically increased hospital supply; [Mommerts and Truskinovsky \(2023\)](#) documented that Medicaid payment rates affect nursing home quality. My results reveal the mirror image: when input costs rise without revenue adjustment, the provider margin contracts. The rate-cost squeeze I document is analogous to the underpayment problem identified by [Grabowski \(2004\)](#) for nursing homes, but operates through the minimum wage rather than reimbursement policy.

The policy implications are significant. Minimum wage increases that are not coordinated with Medicaid rate adjustments create a cost-revenue squeeze that reduces HCBS provider supply, potentially limiting access for low-income beneficiaries who depend on home care services. This represents an unintended tension between two policies with complementary goals: the minimum wage aims to improve the lives of low-wage workers, while Medicaid HCBS aims to serve the most vulnerable populations. States that raise minimum wages should consider concurrent Medicaid rate adjustments to prevent provider market disruption—a recommendation consistent with the ARPA-funded HCBS rate increases that several states implemented during 2021–2023.

8.1 Limitations

Several limitations deserve candid acknowledgment. First, the primary CS estimate is marginally significant ($p = 0.069$). With 9 treated states in 4 cohorts, the design has limited statistical power. Asymptotic cluster-robust standard errors may over-reject in small-cluster settings (Cameron et al., 2008); future work should apply wild cluster bootstrap or randomization inference to sharpen inference.

Second, the American Rescue Plan Act (ARPA) Section 9817 provided a temporary 10-percentage-point increase in the federal Medicaid matching rate for HCBS, beginning in April 2021. Nearly all states used these funds for provider rate increases and workforce bonuses. Because the ARPA HCBS funds overlap with the 2021–2024 post-treatment window, they could attenuate or confound the estimated minimum wage effects. If ARPA funds cushioned the cost-push channel, the true effect of minimum wage increases absent fiscal relief could be larger than what I estimate. A pre-2021 subsample analysis—isolating the 2019 and 2020 cohorts from the ARPA window—would provide a cleaner test, though at the cost of further reducing the treated sample.

Third, the outcome variable (count of billing NPIs) captures provider market participation but not direct employment or service hours. Provider exit from billing does not necessarily mean worker displacement—workers may move to remaining providers—and provider entry may reflect organizational restructuring rather than genuine capacity expansion. Alternative outcome measures such as total beneficiaries served or claims volume per provider would strengthen the analysis.

The welfare calculus is complex. A 15% decline in HCBS provider supply could reduce access to home care for vulnerable beneficiaries, potentially forcing some into more expensive institutional settings—undermining the decades-long rebalancing effort. At the same time, the workers who remain in HCBS earn higher wages, which improves their welfare. The optimal policy likely involves coordinated rate-and-wage increases: raising both the minimum wage and Medicaid HCBS reimbursement rates simultaneously. Several states pursued exactly this approach during the ARPA period, and evaluating the relative effectiveness of coupled versus uncoupled policies is a natural direction for future work.

9. Conclusion

This paper provides the first evidence on how minimum wage policy affects Medicaid home care provider markets, exploiting newly released provider-level T-MSIS data covering all states from 2018 to 2024. The staggered adoption of state minimum wage increases—with 9 states treated across 4 cohorts (2019–2023) and 21 never-treated controls in the primary

Callaway–Sant’Anna design—provides credible variation for a heterogeneity-robust difference-in-differences estimation.

The findings illuminate a tension at the heart of two critical policy domains. Minimum wage increases aim to improve the lives of low-wage workers, many of whom are the same direct care workers who deliver Medicaid-funded home care services. But when these higher labor costs meet the rigid structure of Medicaid reimbursement, the consequences for service delivery capacity become a first-order concern.

Three caveats deserve emphasis. First, T-MSIS measures provider billing activity, not direct worker employment; provider exit from billing does not necessarily mean worker displacement. Second, the analysis captures the total effect of minimum wage increases, including any endogenous Medicaid rate responses; isolating the pure cost-push effect requires data on contemporaneous rate changes that are not systematically available. Third, the comparison between minimum-wage-increasing and federal-minimum states involves geographic and political heterogeneity that, despite robustness checks, may violate parallel trends assumptions.

These findings suggest that the interaction between labor market policy and healthcare delivery infrastructure deserves more attention from both researchers and policymakers. As states continue to raise minimum wages and the federal minimum wage debate persists, understanding the downstream effects on publicly financed healthcare markets is essential for designing policies that achieve their goals without unintended harm to vulnerable populations.

Acknowledgements

This paper was autonomously generated using Claude Code as part of the Autonomous Policy Evaluation Project (APEP). The T-MSIS Medicaid Provider Spending data were obtained from HHS Open Data. State minimum wage data are from the Federal Reserve Economic Data (FRED) system. Population data are from the U.S. Census Bureau’s American Community Survey.

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

Contributors: @olafdrw

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

References

- Baker, Andrew, Bryce Callaway, Scott Cunningham, Andrew Goodman-Bacon, and Pedro HC Sant’Anna**, “Difference-in-Differences Designs: A Practitioner’s Guide,” *arXiv preprint arXiv:2503.13323*, 2025.
- Bureau of Labor Statistics**, “Occupational Outlook Handbook: Home Health and Personal Care Aides,” Technical Report, U.S. Department of Labor 2024.
- Callaway, Bryce and Pedro HC 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.
- Card, David and Alan B Krueger**, “Minimum wages and employment: A case study of the fast-food industry in New Jersey and Pennsylvania,” *American Economic Review*, 1994, *84* (4), 772–793.
- Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer**, “The effect of minimum wages on low-wage jobs,” *Quarterly Journal of Economics*, 2019, *134* (3), 1405–1454.
- Clemens, Jeffrey and Joshua D Gottlieb**, “Do physicians’ financial incentives affect medical treatment and patient health?,” *American Economic Review*, 2014, *104* (4), 1320–1349.
- Dube, Arindrajit**, *Impacts of minimum wages: Review of the international evidence*, HM Treasury, 2019.
- Dustmann, Christian, Attila Lindner, Uta Schönberg, Matthias Umkehrer, and Philipp vom Berge**, “Reallocation effects of the minimum wage,” *Quarterly Journal of Economics*, 2022, *137* (1), 267–328.
- Finkelstein, Amy**, “The effect of public insurance eligibility expansions on provider participation in the program,” *Journal of Public Economics*, 2007, *91* (7-8), 1365–1383.
- 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* (Suppl1), W4–363.
- MACPAC**, “Medicaid and the health care workforce,” Technical Report, Medicaid and CHIP Payment and Access Commission 2023.
- Manning, Alan**, “The elusive employment effect of the minimum wage,” *Journal of Economic Perspectives*, 2021, *35* (1), 3–26.
- Mommaerts, Corina and Yulya Truskinovsky**, “The returns to Medicaid provider payment increases,” *NBER Working Paper*, 2023, (31199).
- Musumeci, MaryBeth and Priya Chidambaram**, “Key state policy choices about Medicaid home and community-based services,” *Kaiser Family Foundation Issue Brief*, 2022.
- PHI**, “Direct Care Workers in the United States: Key Facts,” Technical Report, PHI National 2024.
- Phibbs, Ciaran S et al.**, “The nursing home workforce crisis,” Technical Report, Brookings Institution 2023.
- Rambachan, Ashesh and Jonathan Roth**, “A more credible approach to parallel trends,” *Review of Economic Studies*, 2023, *90* (5), 2555–2591.
- Roth, Jonathan, Pedro HC Sant’Anna, Alyssa Bilinski, and John Poe**, “What’s trending in difference-in-differences? A synthesis of the recent econometrics literature,” *Journal of Econometrics*, 2023, *235* (2), 2218–2244.
- Scales, Kezia and Michael J Lepore**, “The direct care workforce: A comprehensive look at growth and demographics,” *PHI National*, 2023.
- 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 Medicaid Provider Spending

The T-MSIS Medicaid Provider Spending file is derived from the Transformed Medicaid Statistical Information System, the primary federal data source for Medicaid program operations. The file was released by the Department of Health and Human Services on February 9, 2026, through the HHS Open Data portal.

The dataset schema consists of seven fields: billing provider NPI, servicing provider NPI, HCPCS code, claim month, total unique beneficiaries, total claims, and total paid amount. The primary key is the combination of billing NPI, servicing NPI, HCPCS code, and month. Rows with fewer than 12 total claims are suppressed for privacy.

The file covers all 50 states, the District of Columbia, and U.S. territories from January 2018 through December 2024, encompassing both fee-for-service and managed care encounter data. Total coverage: 227,083,361 rows, 617,503 unique billing NPIs, \$1.09 trillion in cumulative payments.

A.2 HCPCS Code Classification

I classify HCPCS codes into HCBS (Medicaid-specific) and non-HCBS categories based on their prefix:

- **HCBS:** T-codes (state HCBS services, e.g., T1019 = personal care/15 min), H-codes (behavioral health, e.g., H2016 = community support/diem), S-codes (temporary state codes, e.g., S5125 = attendant care/15 min)
- **Non-HCBS:** Numeric CPT codes (physician services, e.g., 99213 = office visit), J-codes (drugs), A-codes (ambulance/DME), G-codes (CMS procedures), E-codes (DME), L-codes (orthotics)

T, H, and S codes account for 52% of total T-MSIS spending. The top three codes by cumulative spending—T1019 (\$145.5B), T2016 (\$96.3B), and S5125 (\$62.5B)—are all HCBS services delivered by direct care workers.

A.3 NPPES Linkage

The National Plan and Provider Enumeration System provides the registry of all NPI holders. I extract 21 fields from the bulk CSV, including practice state, ZIP code, entity type, taxonomy codes, and lifecycle dates. The match rate between T-MSIS billing NPIs and NPPES is 99.5%.

A.4 State Minimum Wage Panel

State minimum wage levels are from FRED series STTMINWG[STATE], which provides annual January 1 values. I verify mid-year changes against Department of Labor records. States without state minimum wage laws (AL, LA, MS, SC, TN) are assigned the federal minimum of \$7.25.

Treatment is defined as a state’s first minimum wage increase during 2018–2024. States whose minimum wage increased before 2018 (relative to the 2017 level) and continued increasing are assigned to the cohort of their first increase within the sample period.

B. Identification Appendix

B.1 Treatment Rollout

The Callaway–Sant’Anna estimator uses the following cohorts (states first treated in each year):

- **2019 cohort (4 states):** AR, CT, DE, MA
- **2020 cohort (3 states):** IL, NM, NV
- **2021 cohort (1 state):** VA
- **2023 cohort (1 state):** NE

Twenty-one states remained at the federal minimum (\$7.25) throughout 2018–2024 and serve as the never-treated control group: AL, GA, IA, ID, IN, KS, KY, LA, MS, NC, ND, NH, OK, PA, SC, TN, TX, UT, WI, WV, WY.

An additional 21 states were first treated in 2018 (the first year of the panel) and are excluded from CS estimation due to the absence of a pre-treatment baseline. These states contribute to the supplementary TWFE estimates.

B.2 Event Study Coefficients

B.3 Pre-Trend Testing

I conduct a joint Wald test of the null hypothesis that all pre-treatment event-study coefficients equal zero. The test statistic and p-value are reported in [Table 6](#).

Table 6: Event Study Coefficients: Log(HCBS Providers)

Event Time	ATT	SE
$t - 5$	-0.0925	(0.1068)
$t - 4$	-0.0986	(0.1050)
$t - 3$	-0.0217	(0.1039)
$t - 2$	0.0905*	(0.0524)
$t - 1$	[Reference]	
$t + 0$	-0.0462	(0.0369)
$t + 1$	-0.1417*	(0.0727)
$t + 2$	-0.1784*	(0.1032)
$t + 3$	-0.1849*	(0.1063)
$t + 4$	-0.2230*	(0.1138)
$t + 5$	-0.2664**	(0.1196)
Pre-trend Wald p -value	0.324	

Notes: Callaway–Sant’Anna dynamic aggregation. Reference period: $t = -1$.

C. Robustness Appendix

C.1 Sun–Abraham Comparison

As an alternative to Callaway–Sant’Anna, I estimate a Sun–Abraham interaction-weighted event study using the `sunab()` function in the `fixest` R package (Sun and Abraham, 2021). The Sun–Abraham estimator produces interaction-weighted coefficients for each event time that are robust to treatment effect heterogeneity, using a different aggregation scheme than CS. Post-treatment coefficients are uniformly negative: -0.050 at $e = 0$, growing to -0.253 at $e = 5$, with individual significance at the 1% level for $e = 1, 4, 5$. This confirms the growing provider exit pattern observed in the CS event study (Figure 1).

C.2 Goodman-Bacon Decomposition

The Callaway–Sant’Anna estimator avoids the negative weighting problems that plague TWFE in staggered settings. Nonetheless, the TWFE coefficient of -0.159 is nearly identical to the CS ATT of -0.161 , suggesting that “forbidden comparisons” (already-treated vs. later-treated) receive minimal weight in this setting—consistent with the large share of never-treated states (21 of 51) providing clean control variation.

C.3 Robustness Summary

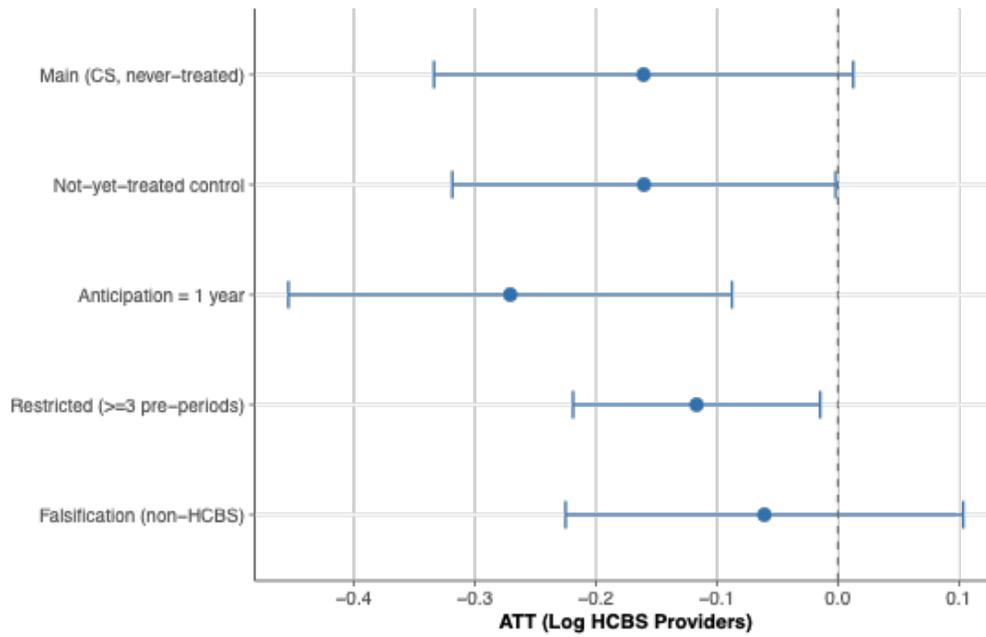


Figure 9: Robustness of Main Result Across Specifications

Notes: Point estimates and 95% confidence intervals for the effect of minimum wage increases on log HCBS providers across CS-based specifications. All estimates use the same outcome variable (log HCBS providers) and comparable ATT estimands. The falsification test uses non-HCBS providers as the outcome.