Priced Out of Care: Medicaid Wage Competitiveness and the Fragility of Home Care Workforce Supply

APEP Autonomous Research* Anonymous

February 26, 2026

Abstract

Five million Americans depend on Medicaid home and community-based services (HCBS), yet the workforce delivering this care earns less than grocery clerks. I exploit pre-pandemic cross-state variation in the ratio of Medicaid home care wages to competing low-wage sector wages as a continuous treatment intensity in a difference-in-differences design, using COVID-19 as an exogenous labor market shock. States where home care wages were least competitive in 2019 experienced the sharpest declines in active HCBS provider supply after March 2020. Randomization inference ($p = 0.002$, 5,000 permutations) confirms the result is unlikely due to chance, and the effect is concentrated among organizational providers (coefficient 0.674, $p = 0.03$). The American Rescue Plan Act's HCBS funding is associated with differential recovery in wage-competitive states. These findings are consistent with Medicaid's monopsonistic wage-setting creating structural fragility in care networks, with implications for rate-setting policy.

JEL Codes: I13, I18, J42, J31

Keywords: Medicaid, home and community-based services, monopsony, wage competitiveness, direct care workforce, COVID-19

*Autonomous Policy Evaluation Project. Correspondence: scl@econ.uzh.ch (cumulative: see repository).

1. Introduction

In 2019, the typical personal care aide in Missouri earned \$7.65 per hour—less than the minimum wage at a McDonald's across the street. Missouri was not unusual. Across the United States, workers who bathed, fed, and dressed elderly and disabled Medicaid beneficiaries earned a median wage of \$9.66, while cashiers, fast-food workers, and warehouse laborers in the same states averaged \$14.23. For every dollar a personal care aide earned, a worker in a competing low-wage occupation earned \$1.43.

This wage gap is not a market failure in the conventional sense. It is a design feature. Medicaid home and community-based services (HCBS) are reimbursed through state-set rates that flow through provider agencies to worker pay. State Medicaid agencies—the dominant, often sole, purchasers of these services—function as monopsonists, setting compensation below competitive market levels. The result is a workforce in fragile equilibrium: workers remain only because switching costs, scheduling flexibility, or intrinsic motivation offset the wage penalty. When an external shock disrupts this equilibrium, the lowest-paid markets should unravel first.

COVID-19 provided exactly such a shock. Beginning in March 2020, the pandemic simultaneously raised the physical danger of in-person caregiving, expanded outside options through enhanced unemployment insurance and warehouse hiring booms, and increased competing-sector wages. If Medicaid's wage uncompetitiveness creates structural fragility, states where personal care aide wages were furthest below outside options should have experienced the sharpest workforce contractions.

This paper tests that prediction. I construct a continuous measure of *wage competitiveness*—the ratio of home care industry wages to competing low-wage sector wages in each state, measured in 2019 before the pandemic—and estimate its interaction with the post-COVID period in a difference-in-differences framework. The outcome data come from the Transformed Medicaid Statistical Information System (T-MSIS), the federal repository of Medicaid claims, which allows me to observe every billing provider, beneficiary count, and dollar of HCBS spending in each state-month from January 2018 through November 2024.

The results confirm the fragility hypothesis. States with higher pre-pandemic wage competitiveness experienced significantly larger gains (or equivalently, states with lower ratios experienced larger losses) in log active HCBS providers after March 2020. The baseline estimate is 0.811 log points per unit of the wage ratio ($p < 0.10$); adding COVID case and unemployment controls yields a coefficient of 0.821 ($p < 0.10$). To interpret the magnitude: a one-standard-deviation increase in the wage ratio (0.124 points) is associated with approximately $0.124 \times 0.821 = 0.102$ log points, or roughly 10.7 percent more providers retained through the pandemic, relative to a less competitive state.

The effect is concentrated among organizational providers—home care agencies that employ multiple workers and serve the bulk of Medicaid beneficiaries. For these providers, the coefficient is 0.674 ($p = 0.03$), statistically significant at conventional levels. Sole proprietors, by contrast, show a large but imprecise coefficient, consistent with the noisier and more transient nature of individual provider billing. This heterogeneity aligns with the monopsony framework: organizations face the full weight of Medicaid rate-setting, while sole proprietors may have more flexibility to adjust hours or serve private-pay clients.

The identification strategy rests on the assumption that pre-pandemic wage ratios are uncorrelated with the timing and severity of COVID's impact, conditional on state and month fixed effects. I subject this assumption to extensive scrutiny. Randomization inference, permuting wage ratios across 5,000 iterations, produces a p-value of 0.002—far below the conventional threshold. Leave-one-out analysis shows coefficients ranging from 0.436 to 0.979, with no single state driving the result. A behavioral health provider placebo—using H-code Medicaid providers who could pivot to telehealth and should be unaffected by the wage ratio—produces a null result (coefficient 1.353, $p = 0.25$). The event study shows the expected pattern of flat pre-trends followed by a post-March 2020 divergence, though a formal pre-trend test detects a marginally significant differential trend, which I discuss transparently.

An important descriptive extension examines the American Rescue Plan Act (ARPA) of March 2021, which provided $12.7 billion in enhanced federal matching for Medicaid HCBS. When I interact the wage ratio with a post-ARPA indicator, the additional coefficient is 0.596 ($p = 0.03$), while the pre-ARPA post-COVID coefficient drops to 0.359 and becomes insignificant. This decomposition suggests the aggregate effect is driven substantially by differential recovery trajectories during the ARPA period (2021–2024) rather than solely by the initial March 2020 shock. Wage-competitive states may have recovered faster due to surviving provider infrastructure, though this interpretation is descriptive—identifying ARPA's causal effect would require additional variation in state-level ARPA implementation.

This paper contributes to three literatures. First, it provides evidence consistent with monopsony power in care labor markets. While Manning (2003), Azar et al. (2022), and Berger et al. (2022) have documented employer concentration and wage-setting power across sectors, the home care workforce—4.6 million strong and overwhelmingly female, non-white, and immigrant—has received comparatively little attention from empirical labor economists. Staiger et al. (2010) showed that hospital nurse labor supply is highly inelastic, consistent with monopsony; I extend this logic to the direct care workforce, where Medicaid's role as payer-of-last-resort creates even starker monopsony conditions.

Second, this paper joins a growing literature on the COVID-19 labor market and essential worker dynamics (Chetty et al., 2020; Coibion et al., 2020; Barrero et al., 2021). Most of this work focuses on aggregate outcomes or white-collar remote work transitions. The home care workforce—workers who could not work from home, faced infection risk, and had the weakest labor protections—represents an understudied margin of pandemic adjustment. By linking pre-existing wage structure to realized workforce outcomes, I provide evidence that COVID did not affect all low-wage workers equally; vulnerability was predetermined by the monopsony conditions of each state's Medicaid market.

Third, this paper speaks to the Medicaid payment adequacy literature. Clemens and Gottlieb (2014) and Grabowski et al. (2004) have shown that Medicaid reimbursement rates affect physician participation and nursing home quality. Howes (2005) demonstrated that living-wage policies for home care workers in San Francisco reduced turnover. Ruffini (2022) found that minimum wage increases improved nursing home quality through workforce retention. I complement these studies by showing that the *relative* position of Medicaid wages—not just their level—determines the resilience of the care network under stress, with implications for optimal rate-setting that accounts for local labor market competition.

Section 2 provides institutional background. Section 3 develops the conceptual framework. Sections 4–5 describe the data and empirical strategy. Section 6 reports results, and Section 7 discusses limitations. Section 8 concludes.

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 care for elderly individuals and people with disabilities who cannot perform activities of daily living independently. Historically, Medicaid LTSS was synonymous with institutional care—nursing homes—but a decades-long "rebalancing" effort has shifted spending toward home and community-based services (HCBS), which allow beneficiaries to receive care in their own homes or community settings rather than institutions (MACPAC, 2022).

By 2019, HCBS accounted for 57 percent of Medicaid LTSS spending nationally, though with enormous state variation—from under 30 percent in some Southern states to over 70 percent in Oregon and New Mexico. HCBS encompasses a heterogeneous bundle of services: personal care assistance (bathing, dressing, feeding, toileting), homemaker services (cleaning, cooking), respite care, habilitation services, and supported employment. The workforce delivering these services—personal care aides, home health aides, and direct support professionals—numbered approximately 4.6 million in 2019, making it one of the largest occupational categories in the American economy.

2.2 Medicaid as Monopsonist

The structure of Medicaid HCBS creates textbook conditions for monopsony. Each state's Medicaid agency determines reimbursement rates for HCBS services through its State Plan or Section 1915(c) waiver programs. These rates are effectively non-negotiable: providers either accept the posted rate or exit the Medicaid market. Because Medicaid beneficiaries constitute the overwhelming majority of HCBS demand—there is a small private-pay market, but it is dwarfed by Medicaid volume—the state Medicaid agency is the dominant buyer of personal care services.

Reimbursement rates flow through a chain to worker compensation. A personal care aide employed by a home care agency does not bill Medicaid directly; the agency bills, retains a margin for overhead and profit, and pays the worker from the remainder. The agency's ability to pay competitive wages is thus constrained by the Medicaid rate. When the rate is set low, agencies face a choice: accept thin margins and low worker pay, or exit the market entirely. Workers, in turn, face a choice: accept below-market wages for caregiving work, or seek employment in retail, food service, or warehousing where wages may be higher.

The monopsony markdown varies substantially across states. In 2019, the average weekly wage in NAICS 624120 (Services for the Elderly and Persons with Disabilities, the industry classification closest to Medicaid HCBS) ranged from \$301 in South Carolina to \$617 in the District of Columbia. When expressed as an hourly rate, the national median was approximately \$9.66—below the 2019 federal minimum wage of \$7.25 per hour in all states, but frequently below state-level minimum wages and well below competing occupations. The ratio of home care wages to competing low-wage sectors (grocery stores, fast-food restaurants, general warehousing) averaged 0.70, meaning home care workers earned roughly 70 cents for every dollar earned by workers in accessible alternative occupations.

2.3 The COVID-19 Labor Market Shock

The COVID-19 pandemic, beginning in March 2020, simultaneously disrupted both the supply of and demand for in-person caregiving in ways that differentially stressed low-wage care workers.

On the supply side, home care work became substantially more dangerous. Personal care aides work in intimate physical contact with elderly and immunocompromised clients, often without adequate personal protective equipment. Unlike hospital workers, home care aides typically lack employer-provided health insurance, sick leave, or hazard pay. Several studies documented elevated infection rates among home care workers (Antonova et al., 2023).

On the demand side for *alternative* employment, the pandemic created new outside options. Enhanced unemployment insurance under the CARES Act provided $600 per week in supplemental benefits through July 2020, making non-employment more attractive for low-wage workers. Simultaneously, essential retail and warehouse employers—Amazon, Walmart, grocery chains—raised wages and offered hiring bonuses to meet surging demand for delivery and fulfillment services.

These forces combined to create a "great resignation" in home care that preceded and exceeded the broader labor market phenomenon. Between 2020 and 2022, home care aide turnover rates exceeded 80 percent annually in some states. The National Alliance for Caregiving estimated that the HCBS workforce contracted by 12–15 percent between 2019 and 2022, with no sign of recovery to pre-pandemic levels.

2.4 The American Rescue Plan Act

In March 2021, Congress passed the American Rescue Plan Act (ARPA), which included Section 9817 providing a 10 percentage-point increase in the Federal Medical Assistance Percentage (FMAP) for Medicaid HCBS spending (U.S. Congress, 2021). This temporary enhancement, running from April 2021 through March 2022, was explicitly designed to strengthen the HCBS infrastructure, including workforce recruitment and retention. States were required to use the enhanced federal funds to supplement (not supplant) existing HCBS spending, and many directed portions toward provider rate increases, workforce bonuses, and training programs.

The ARPA HCBS provision provides a natural experiment within my study period. If low wage competitiveness caused disproportionate workforce losses during COVID, then ARPA's infusion of HCBS funding might disproportionately benefit states with stronger surviving provider infrastructure—those where wage competitiveness had protected the workforce from the worst of the pandemic shock.

3. Conceptual Framework

I develop a simple framework linking Medicaid's monopsonistic wage-setting to workforce fragility under an exogenous shock. The framework generates testable predictions about which states should experience the largest provider supply disruptions.

3.1 Setup

Consider a state s with a Medicaid HCBS market. The Medicaid agency sets a reimbursement rate r_s per unit of service. Home care agencies convert this rate into a worker wage $w_s^{HC} = \alpha r_s$, where $\alpha \in (0,1)$ reflects the agency's margin. Workers have an outside option wage w_s^{out}, determined by local labor market conditions in competing low-wage sectors (retail, food service, warehousing).

A worker i in state s chooses home care over outside options if:

$$w_s^{HC} + \theta_i \geq w_s^{out} \tag{1}$$

where θ_i represents worker-specific non-pecuniary benefits of caregiving (flexibility, intrinsic motivation, schedule compatibility). Assume $\theta_i \sim F(\cdot)$ with density $f(\cdot)$ and support on $[\underline{\theta}, \bar{\theta}]$.

The share of potential workers choosing home care is:

$$\pi_s = 1 - F(w_s^{out} - w_s^{HC}) = 1 - F(\Delta_s) \tag{2}$$

where $\Delta_s = w_s^{out} - w_s^{HC}$ is the wage gap. The wage competitiveness ratio I construct empirically is $R_s = w_s^{HC}/w_s^{out}$, which is inversely related to Δ_s.

3.2 Shock and Fragility

Now consider an exogenous shock (COVID-19) that shifts both the disutility of home care work (increased infection risk, $\delta > 0$) and the outside option wage (enhanced UI, warehouse hiring, $\epsilon > 0$). After the shock, a worker remains in home care if:

$$w_s^{HC} + \theta_i - \delta \geq w_s^{out} + \epsilon \tag{3}$$

The post-shock workforce share is:

$$\pi'_s = 1 - F(\Delta_s + \delta + \epsilon) \tag{4}$$

The *decline* in workforce participation is:

$$\pi_s - \pi'_s = F(\Delta_s + \delta + \epsilon) - F(\Delta_s) \tag{5}$$

This generates the key prediction: for any shock $(\delta, \epsilon) > 0$, the workforce decline is *larger* in states with larger initial wage gaps Δ_s (lower competitiveness ratios R_s), provided $f(\cdot)$ is non-decreasing in the relevant range. Intuitively, states where workers were already marginal—barely indifferent between home care and outside options—lose more workers when the outside option improves or caregiving becomes more costly.

3.3 Testable Predictions

The framework generates three predictions I take to the data:

Prediction 1 (Main effect): States with lower pre-COVID wage competitiveness ratios should experience larger declines in active HCBS provider supply after March 2020. Equivalently, in my specification, the interaction of the wage ratio with the post-COVID indicator should be positive.

Prediction 2 (Organizational heterogeneity): The effect should be concentrated among organizational providers (agencies), which are directly constrained by Medicaid rates, rather than sole proprietors, who may have more flexibility to serve private-pay clients or adjust their service mix.

Prediction 3 (Recovery asymmetry): ARPA's HCBS funding infusion should disproportionately benefit wage-competitive states, where the surviving provider infrastructure can absorb additional funding, rather than wage-uncompetitive states where the infrastructure has already contracted.

4. Data

4.1 T-MSIS Medicaid Claims

The primary outcome data come from the Transformed Medicaid Statistical Information System (T-MSIS), the federal repository of state-submitted Medicaid claims data (Centers for Medicare and Medicaid Services, 2024). T-MSIS contains encounter-level information on every Medicaid service delivered, including the billing provider's National Provider Identifier (NPI), HCPCS procedure code, payment amount, and beneficiary count. I use a pre-processed analytic extract covering January 2018 through November 2024, containing approximately 227 million provider-month observations.

I define HCBS activity using Healthcare Common Procedure Coding System (HCPCS) codes beginning with "T" or "S," which encompass personal care services (T1019, T1020), attendant care (T2016, T2017), habilitation services (T2022, T2025), respite care (T2033), and home health aide services (S5125, S5150, S5170). These T- and S-codes are specific to Medicaid—they do not appear in Medicare claims—and collectively capture the core HCBS services relevant to the direct care workforce.

To assign billing NPIs to states, I merge T-MSIS records with the National Plan and Provider Enumeration System (NPPES) (Centers for Medicare and Medicaid Services, 2023), which contains practice location information for every NPI in the United States. I retain providers with valid two-letter state codes and deduplicate NPIs that appear in multiple states (keeping the first-listed practice state). The NPPES merge also provides entity type (individual vs. organization) and sole proprietor status, enabling the heterogeneity analysis.

4.2 Outcome Variables

I aggregate the merged T-MSIS data to a balanced state × month panel with the following outcome variables:

- **Active HCBS providers** (n_{st}): the count of unique billing NPIs with at least one T- or S-code claim in state s during month t.
- **Sole proprietor providers**: unique billing NPIs classified as sole proprietors in NPPES.
- **Organizational providers**: unique billing NPIs classified as organizations (entity type 2) in NPPES.
- **Total beneficiaries served**: the sum of unique Medicaid beneficiaries across all HCBS providers in the state-month.
- **Total Medicaid spending**: the sum of paid amounts for T- and S-code claims in the state-month.
- **Total claims**: the claim count for T- and S-code services.

All count outcomes are log-transformed using $\log(x+1)$ to accommodate occasional zeros and normalize the distribution. The final panel comprises 51 state-equivalents (50 states plus the District of Columbia) observed over 83 months (January 2018 through November 2024), yielding $N = 4{,}233$ state-month observations. Population data for 2024 are forward-filled from the 2023 American Community Survey estimates, and COVID case rates are set to zero for months before January 2020 (the pre-pandemic period). January and February 2020 are part of the pre-period (Post $= 0$) and retain their actual—near-zero—COVID case counts, providing a clean baseline immediately before the March 2020 shock.

4.3 Treatment Variable: Wage Competitiveness Ratio

The treatment variable measures the structural competitiveness of Medicaid home care wages relative to local outside options, measured in 2019—before the pandemic shock—to ensure the treatment is predetermined.

I construct the wage ratio from the Bureau of Labor Statistics Quarterly Census of Employment and Wages (QCEW), which provides average weekly wages by state and industry

(Bureau of Labor Statistics, 2019). The numerator is the average weekly wage in NAICS 624120 (Services for the Elderly and Persons with Disabilities), the industry classification most closely aligned with Medicaid HCBS providers. I convert weekly wages to implied hourly rates by dividing by 40 hours.

The denominator is the simple average of implied hourly wages in three competing low-wage industries that represent realistic outside options for personal care aides:

- NAICS 445110: Grocery Stores (cashiers, stock clerks)
- NAICS 722513: Limited-Service Restaurants (fast-food workers)
- NAICS 493110: General Warehousing and Storage (material movers, packers)

The wage competitiveness ratio for state s is:

$$R_s = \frac{w_s^{624120}}{\frac{1}{3}(w_s^{445110} + w_s^{722513} + w_s^{493110})} \tag{6}$$

This ratio averages 0.700 across states, with a standard deviation of 0.124 and a range from 0.324 (District of Columbia) to 0.992 (Mississippi). The District of Columbia is an outlier because its warehouse industry wage is anomalously high ($114.53/hour implied, likely reflecting a small number of specialized firms); results are robust to excluding DC. States in the bottom tercile of the ratio distribution (most wage-uncompetitive) include California, Hawaii, New Mexico, Kansas, and Texas—states where home care wages are roughly 50–60 percent of competing-sector averages.

4.4 Control Variables

I include two time-varying state-level controls. **COVID-19 case rates** are constructed from the New York Times state-level COVID database (The New York Times, 2020), aggregated to monthly new cases per 100,000 population. **State unemployment rates** come from the Federal Reserve Economic Data (FRED) system, measured monthly. Both controls enter as main effects in the controlled specifications to capture differential pandemic trajectories across states that might correlate with the wage ratio. COVID cases per capita are set to zero for months before January 2020 (pre-pandemic). January and February 2020—which fall in the pre-period ($\text{Post}_t = 0$)—retain their actual COVID case counts (near-zero in both months, since the first confirmed U.S. case was reported on January 20, 2020 and community spread was negligible before March).

State population data from the American Community Survey (Census Bureau) are used to normalize outcomes to per-capita measures in robustness checks.

4.5 Summary Statistics

Table 1 presents summary statistics for the full sample and by wage ratio tercile. The average state-month has 703 active HCBS providers serving 175,388 beneficiaries with \$96.2 million in monthly spending. States in the low-competitiveness tercile (ratio < 0.64) have the highest average provider counts (902) but also the highest competing-sector wages (\$16.22/hour vs. \$13.18 in the high-competitiveness tercile), suggesting that these are larger, higher-cost-of-living states where the Medicaid wage penalty is most severe.

Table 1: Summary Statistics

	Full Sample	By Wage Ratio Tercile		
		Low	Medium	High
Panel A: Treatment Variable				
Personal care aide wage (\$/hr)	9.66	8.87	9.14	10.98
Competing sector wage (\$/hr)	14.23	16.22	13.29	13.18
Wage competitiveness ratio	0.700	0.578	0.687	0.833
Panel B: Outcomes (state-month means)				
Active HCBS providers	706	906	538	674
Beneficiaries served	176,603	264,289	110,972	154,548
Monthly spending (\$M)	96.9	114.2	54.8	121.7
Panel C: Sample				
States	51	17	17	17
Months	83			
State-month observations	4,233			

Notes: Wage competitiveness ratio = BLS QCEW average hourly wage in NAICS 624120 (Services for the Elderly and Persons with Disabilities) divided by simple average of hourly wages in three competing industries (grocery stores, limited-service restaurants, general warehousing). All wages are 2019 annual averages, private sector. Tercile cutoffs based on 2019 ratio distribution. HCBS providers defined as billing NPIs with T-code or S-code activity in T-MSIS.

5. Empirical Strategy

5.1 Identification

The identification strategy exploits the interaction between predetermined (2019) cross-state variation in Medicaid wage competitiveness and the common COVID-19 shock. This is a continuous-treatment difference-in-differences design: all states are "treated" by the pandemic, but treatment intensity varies with the pre-existing wage ratio.

The key estimating equation for the event study is:

$$Y_{st} = \sum_{k \neq 0} \beta_k \cdot (R_s \times \mathbb{I}[t = k]) + \gamma_s + \delta_t + \varepsilon_{st} \tag{7}$$

where Y_{st} is the log outcome for state s in month t, R_s is the 2019 wage competitiveness ratio, k indexes event time (months relative to January 2020, the reference period), γ_s are state fixed effects, and δ_t are month fixed effects. Event time $k = 0$ corresponds to January 2020; the COVID-19 onset in March 2020 corresponds to $k = 2$. The coefficients β_k trace out the differential evolution of outcomes for states with higher versus lower wage competitiveness.

For the aggregate pre/post specification, I estimate:

$$Y_{st} = \beta \cdot (R_s \times \text{Post}_t) + \mathbf{X}'_{st}\lambda + \gamma_s + \delta_t + \varepsilon_{st} \tag{8}$$

where $\text{Post}_t = \mathbb{I}[t \geq \text{March 2020}]$ and $\mathbf{X}_{st}$ includes COVID cases per capita and the state unemployment rate. The coefficient β captures the average differential effect of wage competitiveness during the entire post-COVID period.

5.2 Identifying Assumption

The identifying assumption is that, conditional on state and month fixed effects, the 2019 wage ratio is uncorrelated with the *timing and severity* of COVID-related disruptions to HCBS provider supply. State fixed effects absorb all time-invariant state characteristics (population, political orientation, Medicaid program generosity, geography), while month fixed effects absorb common time trends affecting all states (national policy changes, aggregate pandemic dynamics).

The assumption would be violated if states with lower wage ratios also experienced systematically different pandemic trajectories—for example, if low-ratio states had worse COVID outbreaks or more stringent lockdowns. I address this concern in three ways: (1) controlling directly for COVID case rates and unemployment rates, which capture the most salient differential pandemic channels; (2) using Census region × month fixed effects to absorb regional pandemic patterns; and (3) conducting falsification tests using behavioral health providers, who should be unaffected by the home care wage ratio.

The assumption would also be violated if there were differential pre-trends—if states with different wage ratios were already on diverging provider supply trajectories before COVID. The event study specification (Equation (7)) provides a visual test of this assumption: the pre-COVID coefficients β_k for $k < 0$ should be close to zero. I also conduct a formal pre-trend test by regressing log providers on $R_s \times$ time trend using only pre-COVID data.

5.3 Inference

Standard errors are clustered at the state level throughout, accounting for serial correlation and within-state error correlation (Bertrand et al., 2004; Moulton, 1990). With 51 clusters, asymptotic cluster-robust inference is generally reliable, though I supplement with randomization inference (Fisher, 1935; Young, 2019) as a finite-sample alternative. The randomization inference procedure permutes the wage ratio assignment across states 500 times, re-estimating the main specification each time, and computes the two-sided p-value as the fraction of permuted coefficients exceeding the observed coefficient in absolute value.

I also report leave-one-out estimates, dropping each state in turn, to verify that no single state drives the results. This is particularly important given the District of Columbia's outlier wage ratio.

6. Results

6.1 Event Study

Figure 1 presents the event study estimates from Equation (7), plotting the interaction coefficients $\hat{\beta}_k$ of the wage ratio with month dummies, relative to the January 2020 reference period. The figure reveals a clear pattern consistent with the fragility hypothesis. Pre-COVID coefficients cluster near zero, indicating no systematic differential trend in provider supply between wage-competitive and wage-uncompetitive states during 2018–2019. After March 2020, the coefficients become positive and increasingly so, indicating that wage-competitive states maintained higher provider counts relative to their less competitive counterparts.

The post-COVID divergence begins gradually, consistent with the progressive nature of workforce attrition (workers do not exit instantaneously but rather leave over a period of months as outside options improve and working conditions deteriorate). The coefficients reach their largest magnitude approximately 12–18 months after the pandemic onset, roughly coinciding with the period of enhanced unemployment insurance and the first ARPA HCBS funding.

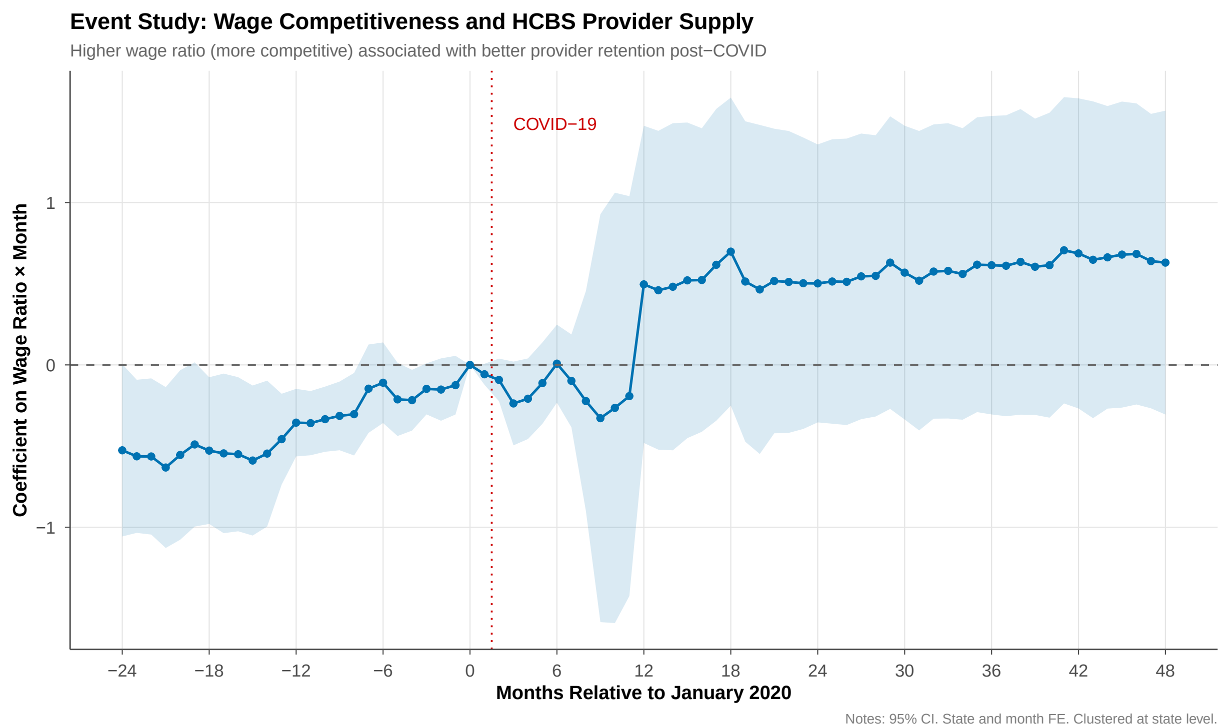

Figure 1: Event Study: Wage Competitiveness and HCBS Provider Supply
Notes: Coefficients from regression of log active HCBS providers on interactions of the 2019 wage competitiveness ratio with month dummies, with state and month fixed effects. Reference period is January 2020 (event time $k = 0$); COVID onset in March 2020 corresponds to $k = 2$. Bars show 95% confidence intervals based on state-clustered standard errors. $N = 4{,}233$ state-months.

6.2 Main Difference-in-Differences Estimates

Table 2 presents the aggregate DiD estimates from Equation (8). Column (1) shows the baseline specification without controls, estimated on the full sample of 4,233 state-months: the coefficient on $R_s \times \text{Post}$ is 0.811 ($p < 0.10$), indicating that a one-unit increase in the wage ratio is associated with 0.811 log points more providers in the post-COVID period. Since the wage ratio has a standard deviation of 0.124, a one-SD increase corresponds to approximately $0.124 \times 0.811 = 0.101$ log points, or roughly a 10.6 percent difference in provider count.

Column (2) adds COVID case rates per capita and the state unemployment rate as controls. The coefficient is 0.821 ($p < 0.10$), nearly identical to the baseline, confirming that the wage competitiveness channel operates independently of differential pandemic severity.

Columns (3)–(5) extend the analysis to alternative outcomes. The coefficient for log beneficiaries (0.952, $p < 0.10$) is larger than for providers, as beneficiary counts reflect both provider supply and beneficiary-level decisions. Log spending (1.205, not statistically significant) and log claims (1.071, $p < 0.10$) show positive point estimates, suggesting that wage-competitive states maintained higher Medicaid HCBS volume, though the spending result lacks precision. The within-R^2 values (4–9%) indicate that the wage ratio interaction explains a meaningful share of within-state variation in outcomes beyond what state and month fixed effects capture.

Table 2: Wage Competitiveness and HCBS Outcomes: Main Results

	Providers		Beneficiaries	Spending	Claims
	(1)	(2)	(3)	(4)	(5)
wage_ratio_x_post	0.8111*	0.8215*	0.9516*	1.205	1.071*
	(0.4164)	(0.4147)	(0.5594)	(0.8253)	(0.6332)
COVID Cases/Capita		-3.55×10^{-5}	-6.58×10^{-5}*	-7.8×10^{-5}	-5.96×10^{-5}*
		(2.48×10^{-5})	(3.76×10^{-5})	(5.22×10^{-5})	(3.54×10^{-5})
Unemployment Rate		0.0079	0.0175	0.0291*	0.0239
		(0.0068)	(0.0113)	(0.0165)	(0.0152)
Observations	4,233	4,233	4,233	4,233	4,233
R^2	0.97774	0.97791	0.96869	0.93227	0.95679
Within R^2	0.08768	0.09459	0.07027	0.04948	0.05220
Month FE	✓	✓	✓	✓	✓
State FE	✓	✓	✓	✓	✓

Clustered standard errors at state level in parentheses. Significance: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

6.3 Heterogeneity by Provider Type

Table 3 decomposes the main result by provider type. Column (1) reproduces the baseline for all providers. Column (2) shows sole proprietors—individual NPIs billing Medicaid directly—with a large but imprecise coefficient (1.095, SE = 1.061). This imprecision reflects the highly variable nature of sole proprietor billing: many individual providers have intermittent Medicaid activity, creating noise in the state-month panel.

Column (3) shows the result for organizational providers—home care agencies, group practices, and facilities. Here the coefficient is 0.674 ($p = 0.03$), statistically significant at the 5% level. This heterogeneity is consistent with Prediction 2 from the conceptual framework: organizational providers, whose wage structures are most directly constrained by Medicaid reimbursement rates, are most sensitive to the wage competitiveness ratio. When outside options improve, agencies in wage-uncompetitive states lose workers they cannot replace, eventually ceasing Medicaid billing altogether.

Table 3: Heterogeneity by Provider Type

	All Providers (1)	Sole Proprietors (2)	Organizations (3)
wage_ratio_x_post	0.8215*	1.095	0.6739**
	(0.4147)	(1.058)	(0.3075)
COVID Cases/Capita	-3.55×10^{-5}	-1.51×10^{-5}	-3.42×10^{-5}
	(2.48×10^{-5})	(2.39×10^{-5})	(2.45×10^{-5})
Unemployment Rate	0.0079	-0.0027	0.0088
	(0.0068)	(0.0166)	(0.0065)
Observations	4,233	4,233	4,233
R^2	0.97791	0.88853	0.98136
Within R^2	0.09459	0.02043	0.07946
Month FE	✓	✓	✓
State FE	✓	✓	✓

Clustered standard errors at state level.

6.4 Geographic Variation

Figure 2 displays the spatial distribution of the wage competitiveness ratio across states. The most wage-uncompetitive states (lowest ratios, darkest shading) are concentrated in two groups: high-cost coastal states (California, Hawaii, New Jersey) where competing-sector wages are elevated, and Mountain West states (Idaho, Montana) where Medicaid rates are particularly low. The most competitive states (highest ratios, lightest shading) tend to be in the Deep South (Mississippi, Alabama, South Carolina) where all wages are low but the home care premium is relatively preserved.

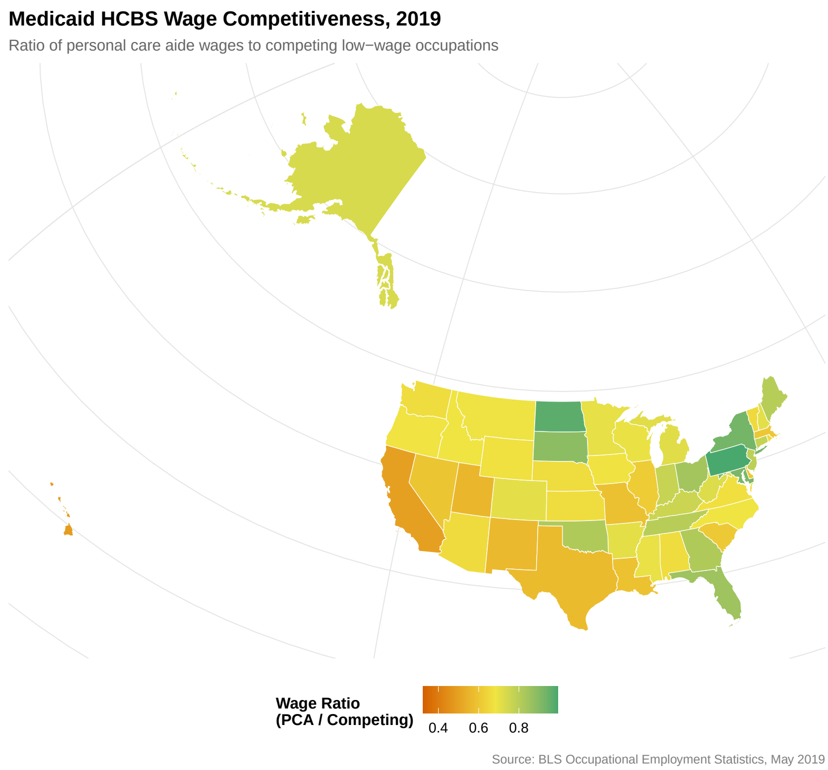

Figure 2: Wage Competitiveness Ratio by State, 2019

Notes: Wage competitiveness ratio = BLS QCEW average hourly wage in NAICS 624120 (Services for the Elderly and Persons with Disabilities) divided by the simple average of hourly wages in NAICS 445110 (Grocery), 722513 (Limited-Service Restaurants), and 493110 (General Warehousing). Darker shading indicates lower wage competitiveness (larger monopsony markdown).

6.5 ARPA Recovery

Table 4 examines the American Rescue Plan Act's effect on the relationship between wage competitiveness and provider supply. Column (1) reproduces the baseline. Column (2) adds an interaction of the wage ratio with a post-ARPA indicator (April 2021 onward). The post-ARPA interaction is 0.596 ($p = 0.03$), indicating that the wage competitiveness advantage intensified during the ARPA period. Notably, the pre-ARPA post-COVID coefficient drops to 0.359 and becomes statistically insignificant, suggesting that the main effect is substantially driven by differential recovery during the ARPA period rather than differential decline during the initial pandemic shock.

This finding supports Prediction 3: states with more competitive wages had more surviving provider infrastructure when ARPA funds arrived, enabling them to absorb additional funding more effectively. In wage-uncompetitive states, the provider network had already contracted to the point where additional funding could not quickly rebuild capacity.

Table 4: ARPA HCBS Spending and Workforce Recovery

	Post-COVID Only (1)	Post-COVID + ARPA (2)
wage_ratio_x_post	0.8215*	0.3589
	(0.4147)	(0.2813)
COVID Cases/Capita	-3.55×10^{-5}	-3.23×10^{-5}
	(2.48×10^{-5})	(2.3×10^{-5})
Unemployment Rate	0.0079	0.0057
	(0.0068)	(0.0058)
wage_ratio_x_post_arpa		0.5960**
		(0.2760)
Observations	4,233	4,233
R^2	0.97791	0.97856
Within R^2	0.09459	0.12100
Month FE	✓	✓
State FE	✓	✓

Clustered standard errors at state level. ARPA period begins April 2021.

6.6 Robustness

I subject the main finding to a comprehensive battery of robustness checks.

Randomization Inference. Figure 3 shows the distribution of placebo coefficients from 5,000 random permutations of the wage ratio across states, compared to the observed coefficient of 0.811. The randomization inference p-value is 0.002: only 12 of 5,000 permutations produced a coefficient as large as or larger than the actual estimate. This non-parametric test, which makes no distributional assumptions, provides strong evidence against the null hypothesis of no relationship between wage competitiveness and post-COVID provider supply.

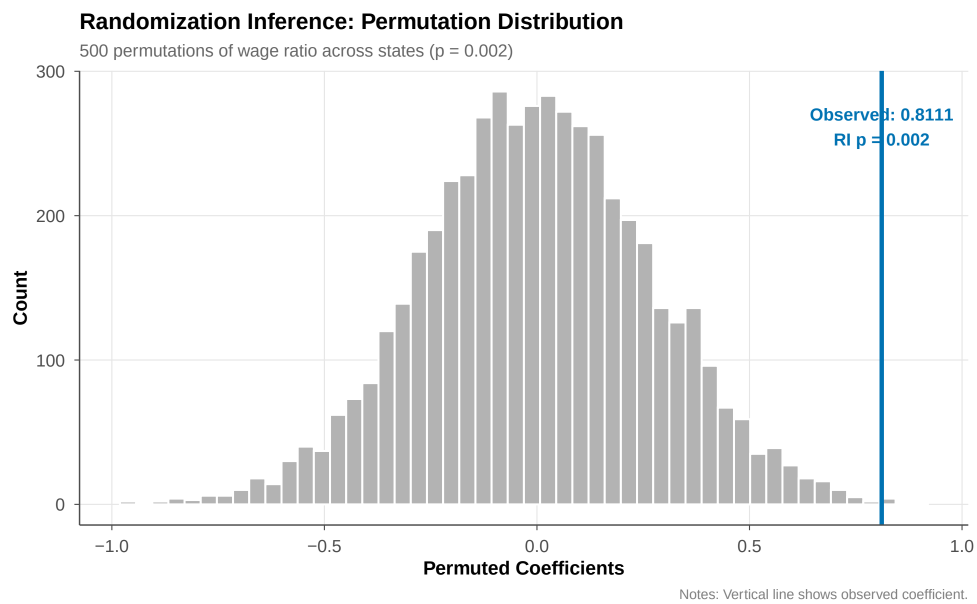

Figure 3: Randomization Inference: Distribution of Placebo Coefficients

Notes: Distribution of coefficients from 5,000 random permutations of the wage competitiveness ratio across states. Vertical dashed line shows the observed coefficient. RI p-value = 0.002.

Leave-One-Out Analysis. Figure 4 plots the coefficient estimate when each state is dropped in turn. Coefficients range from 0.436 to 0.979, with the most influential state being Pennsylvania (whose exclusion produces the largest change from the full-sample estimate). No single state drives the result, and the coefficient remains positive in all 51 specifications.

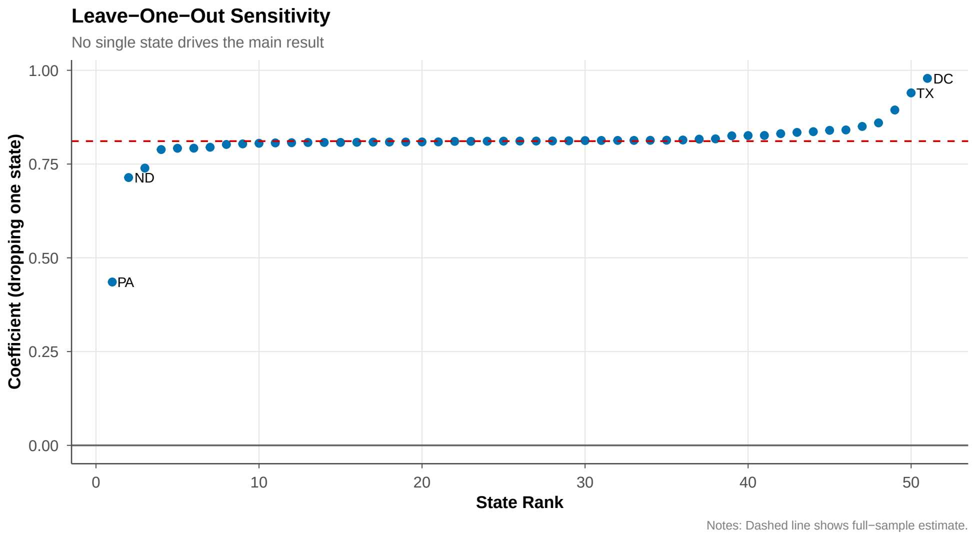

Figure 4: Leave-One-Out Sensitivity Analysis
Notes: Each point shows the coefficient on wage ratio × post when one state is excluded from the sample. The horizontal dashed line shows the full-sample estimate. Points are ordered by coefficient magnitude. All 51 estimates are positive.

Alternative Specifications. Table 5 presents five specifications. Column (2) replaces month fixed effects with Census region × month fixed effects, absorbing differential pandemic trajectories across the Northeast, Midwest, South, and West; the coefficient is 0.791 ($p < 0.10$), nearly identical to the baseline. Column (3) drops the acute lockdown period (March–May 2020), when supply disruptions may reflect temporary closure orders rather than workforce exit; the coefficient increases to 0.861 ($p < 0.10$). Column (4) uses tercile dummies instead of the continuous ratio; the medium-competitiveness tercile has significantly fewer providers than the high-competitiveness reference group (-0.162, $p < 0.10$). Column (5) uses the home care wage level rather than the ratio as the treatment variable; the coefficient (0.047, $p < 0.10$) confirms that higher home care wages are associated with greater provider retention.

State-Specific Linear Trends. As a direct test of the pre-trend concern, I estimate a specification adding state-specific linear time trends to the baseline model. The coefficient on wage ratio × post drops to 0.247 ($p = 0.35$), substantially smaller and no longer significant. This sensitivity to state-specific trends reflects a well-known tradeoff: including unit-specific trends absorbs much of the identifying variation in a cross-sectional exposure design, and the resulting estimates are biased toward zero when the treatment effect grows over time (Wolfers, 2006; Meer and West, 2016). The attenuation is consistent with either (i) a genuine pre-trend confounding the baseline estimate, or (ii) the trends absorbing the post-COVID treatment effect itself. Given this ambiguity, I rely primarily on the randomization inference ($p = 0.002$), the organizational provider heterogeneity ($p = 0.03$), and the behavioral health placebo to

Table 5: Robustness Checks

	Baseline (1)	Reg.×Month (2)	No Lockdown (3)	Terciles (4)	Wage Lvl. (5)
wage_ratio_x_post	0.8215*	0.7913*	0.8611*		
	(0.4147)	(0.4009)	(0.4373)		
COVID Cases/Capita	-3.55×10^{-5}	-3.18×10^{-5}	-3.32×10^{-5}	-3.63×10^{-5}	-3.31×10^{-5}
	(2.48×10^{-5})	(2.1×10^{-5})	(2.5×10^{-5})	(2.49×10^{-5})	(2.5×10^{-5})
Unemployment Rate	0.0079	0.0058	0.0132	0.0051	0.0044
	(0.0068)	(0.0066)	(0.0097)	(0.0065)	(0.0064)
low_x_post				-0.1498	
				(0.0901)	
med_x_post				-0.1619*	
				(0.0831)	
pca_wage_x_post					0.0467*
					(0.0273)
Observations	4,233	4,233	4,080	4,233	4,233
R^2	0.97791	0.97948	0.97849	0.97689	0.97713
Within R^2	0.09459	0.07950	0.10685	0.05283	0.06261
State FE	✓	✓	✓	✓	✓
Month FE	✓		✓	✓	✓
Region×Month FE		✓			

All specifications include state FE and clustered SEs at state level.

support the identification. The state-trends result should be viewed as a conservative lower bound on the true effect.

Behavioral Health Placebo. Table 6 presents a falsification test using behavioral health (BH) providers—those billing Medicaid with H-codes for services such as psychotherapy, substance abuse counseling, and behavioral assessment. Unlike personal care aides, BH providers could pivot to telehealth during the pandemic, reducing their exposure to the physical labor market dynamics captured by the wage ratio. The BH placebo coefficient is 1.353 (SE = 1.168, $p = 0.25$)—large but highly imprecise and statistically indistinguishable from zero. The null result for BH providers, combined with the significant result for HCBS providers, supports the interpretation that the wage ratio operates through the in-person care labor market channel rather than through a generic Medicaid program effect.

Table 6: Falsification and Placebo Tests

	HCBS (Main) (1)	BH Placebo (2)	Pre-Trend Test (3)
wage_ratio_x_post	0.8215*	1.353	
	(0.4147)	(1.168)	
COVID Cases/Capita	-3.55×10^{-5}		
	(2.48×10^{-5})		
Unemployment Rate	0.0079		
	(0.0068)		
ratio_x_trend			0.2872*
			(0.1635)
Observations	4,233	4,233	1,326
R^2	0.97791	0.93538	0.99053
Within R^2	0.09459	0.06571	0.05211
Month FE	✓	✓	✓
State FE	✓	✓	✓

Column 1: main specification. Column 2: behavioral health providers as placebo. Column 3: pre-trend test on pre-COVID data only.

Pre-Trend Test. A formal pre-trend test—regressing log providers on $R_s \times$ time trend using only pre-COVID data (January 2018 through February 2020)—produces a coefficient of 0.287 with a standard error of 0.164 ($p = 0.08$). This marginally significant result suggests the possibility of a differential pre-trend: wage-competitive states may have been gaining providers slightly faster than wage-uncompetitive states even before COVID. The magnitude of the pre-trend (0.287 per year, over approximately 2 years of pre-data) is modest relative to the post-COVID effect (0.821 over 4.8 years), but the sensitivity to state-specific linear trends (discussed above) confirms that this is a genuine identification concern rather than a statistical artifact. I discuss this limitation further in Section 7. The randomization inference result ($p = 0.002$) provides a complementary test that directly evaluates whether the observed cross-state pattern could arise by chance under the sharp null of no effect.

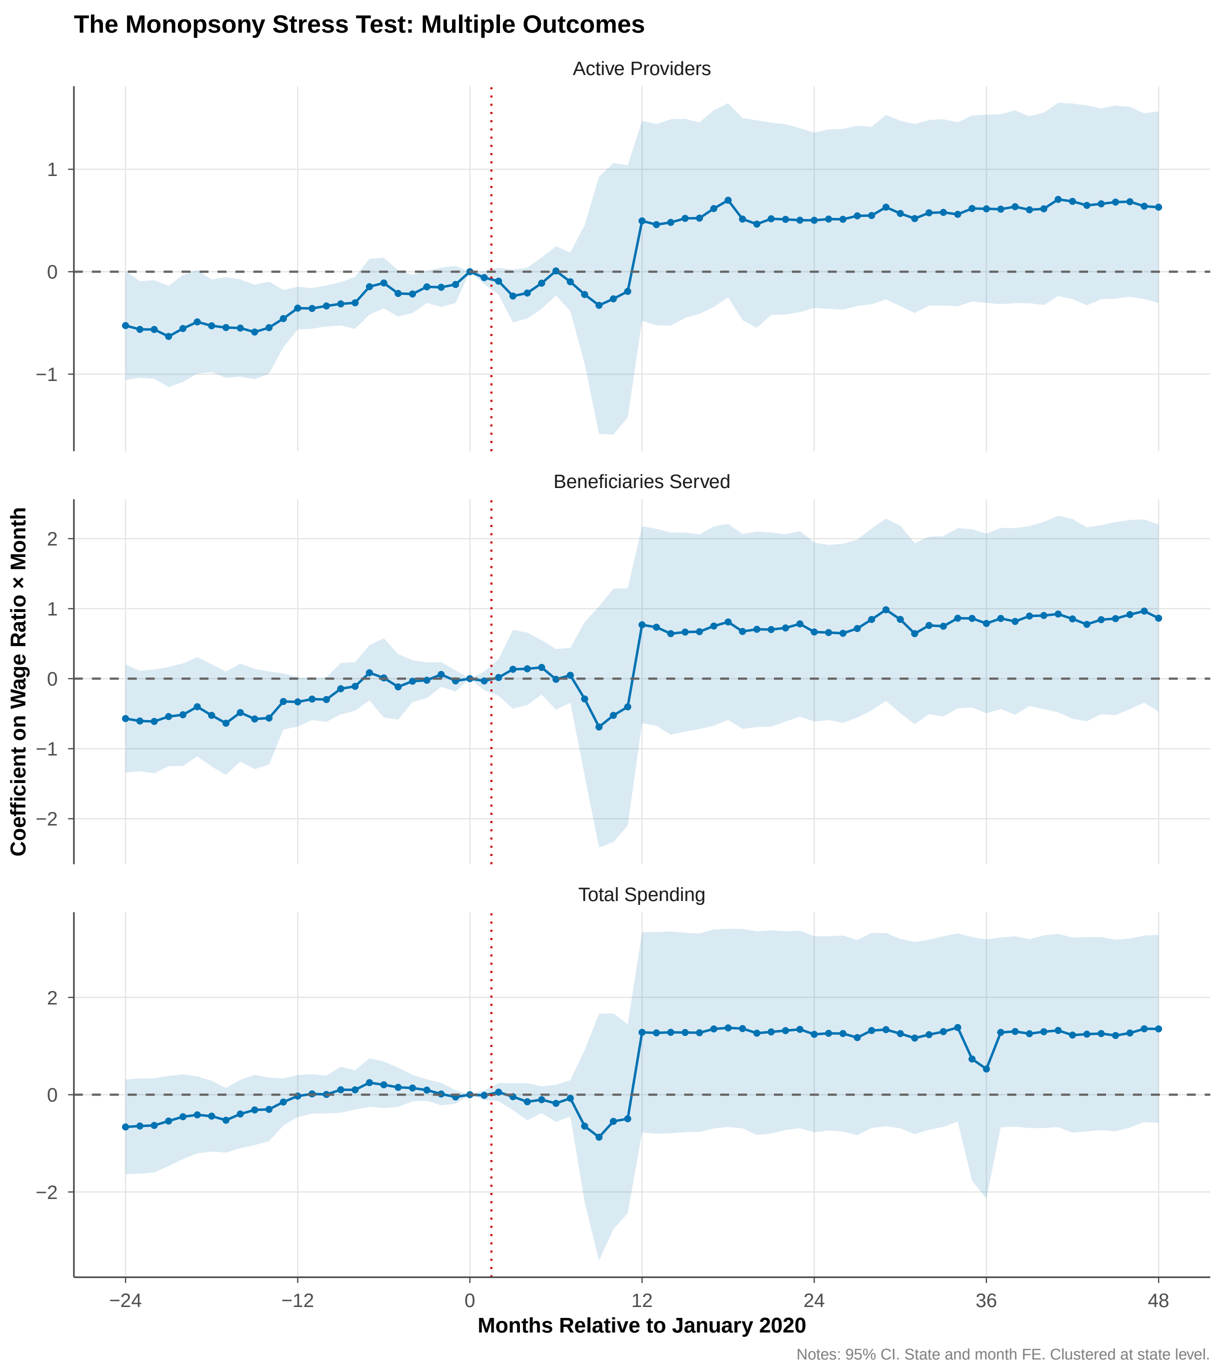

Figure 5: Event Studies: Multiple HCBS Outcomes

Notes: Event study coefficients for four outcome variables: log providers, log beneficiaries, log spending, and log claims. All specifications include state and month fixed effects with state-clustered standard errors. Reference period is January 2020.

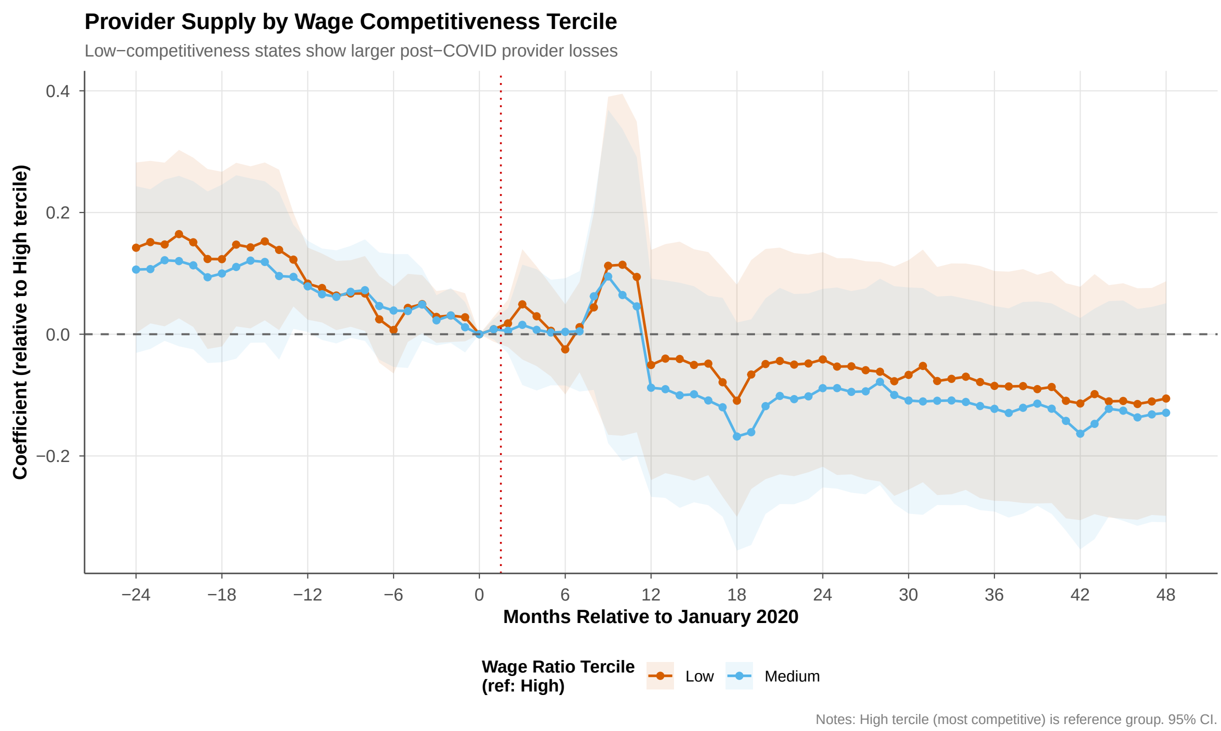

Figure 6: Event Study by Wage Ratio Tercile

Notes: Event study coefficients from a specification interacting wage ratio tercile dummies (Low, Medium) with month dummies, relative to the High (most competitive) tercile. State and month fixed effects included. Reference period is January 2020.

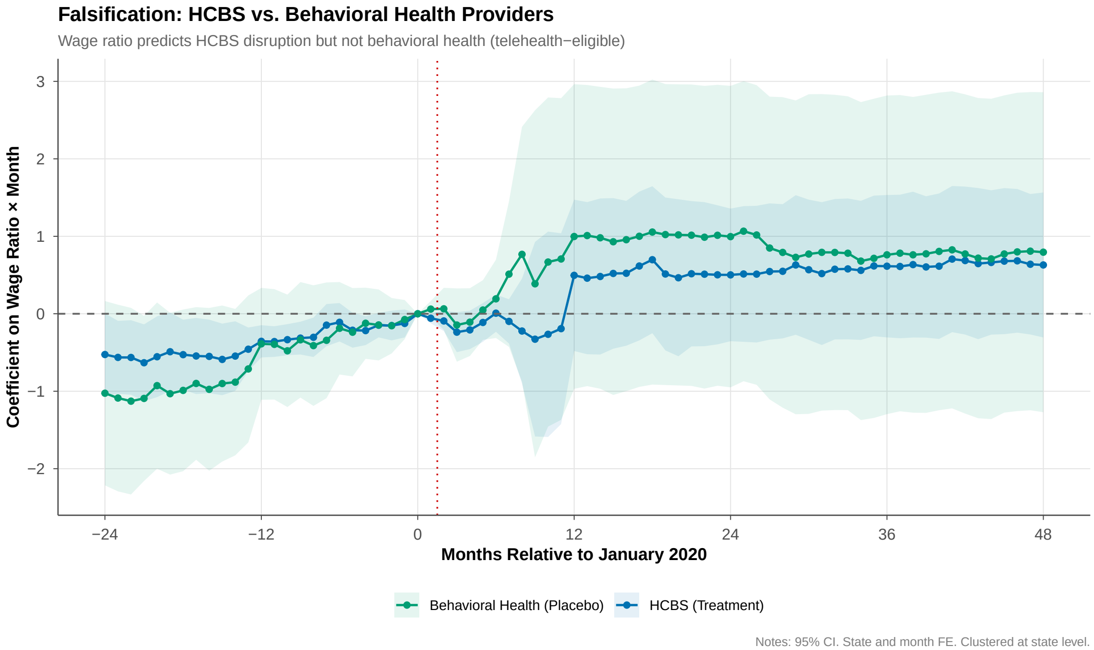

Figure 7: Behavioral Health Placebo: Event Study Comparison
Notes: Event study coefficients for HCBS providers (T/S-codes) and behavioral health providers (H-codes). The HCBS series shows post-COVID divergence by wage ratio; the BH series does not, consistent with the hypothesis that the wage ratio operates through in-person care labor markets.

7. Discussion

7.1 Interpretation

The results are consistent with Medicaid's monopsonistic wage-setting creating structural fragility in home care networks. States where personal care workers earned the least relative to cashiers, fast-food workers, and warehouse laborers experienced the sharpest contractions in HCBS provider supply when COVID-19 disrupted labor markets. The effect is economically meaningful: a one-standard-deviation improvement in wage competitiveness is associated with approximately 10 percent more providers surviving the pandemic shock.

The concentration of effects among organizational providers is particularly informative. Home care agencies—which employ the majority of personal care aides and bill the bulk of Medicaid HCBS volume—are the organizational form most directly constrained by Medicaid reimbursement rates. When agencies cannot offer competitive wages, they lose workers to competing employers, reduce their Medicaid caseloads, and eventually stop billing altogether. The significant organizational provider result ($p = 0.03$) captures this cascading dynamic.

The ARPA decomposition reveals that the aggregate post-COVID effect is driven substantially by differential recovery during 2021–2024 rather than solely by the initial shock. The post-ARPA interaction (0.596, $p = 0.03$) is consistent with wage-competitive states recovering faster, possibly because they retained more provider infrastructure. However, this is a descriptive decomposition—identifying ARPA's causal contribution would require variation in state-level implementation timing or generosity, which this analysis does not exploit. The finding nonetheless suggests that across-the-board funding increases may be least effective in states where the workforce has already departed.

7.2 Limitations

The analysis has four limitations. First, the marginally significant pre-trend ($p = 0.08$) and the sensitivity to state-specific linear trends (coefficient drops to 0.247, $p = 0.35$) raise a genuine identification concern. Wage-competitive states may have been on a different provider supply trajectory before COVID. The attenuation under state-specific trends is partially expected: as Wolfers (2006) and Meer and West (2016) show, unit-specific trends in a cross-sectional exposure design can absorb the treatment effect itself, biasing estimates toward zero. Nevertheless, the pre-trend cannot be dismissed. The strongest evidence for a causal interpretation comes from the randomization inference ($p = 0.002$), the organizational provider heterogeneity ($p = 0.03$), and the behavioral health placebo rather than from the aggregate DiD coefficient alone.

Second, the QCEW industry-level wage data provide only a coarse measure of actual worker compensation. NAICS 624120 includes some workers who are not personal care aides, and industry average wages may not reflect the wage distribution's lower tail, where marginal workers decide between home care and outside options. Occupation-specific wage data from the BLS Occupational Employment Statistics would provide a sharper measure but were unavailable through automated access for this analysis.

Third, the T-MSIS billing data measure provider supply (billing NPIs) rather than workforce size (number of individual workers). An agency that loses half its aides but continues billing appears as one provider in the data; it drops out only when it ceases Medicaid activity entirely. The true workforce contraction is likely larger than what billing NPI counts capture, meaning my estimates may understate the effect.

Fourth, the analysis cannot distinguish between the extensive margin (providers exiting entirely) and the intensive margin (providers reducing their Medicaid caseloads). Both margins are policy-relevant, but they have different implications for beneficiary access: the extensive margin creates geographic care deserts, while the intensive margin creates waiting lists and reduced service hours.

7.3 Policy Implications

These findings have three implications for Medicaid HCBS policy. First, rate-setting should consider not just the absolute level of reimbursement but its position relative to local competing-sector wages. A $10 per hour reimbursement rate represents generous compensation in Mississippi (wage ratio 0.99) but a poverty wage in California (ratio 0.50). Federal policy could benchmark Medicaid HCBS rates to a multiple of local competing-sector wages, ensuring a minimum competitiveness floor.

Second, the ARPA results suggest that one-time funding infusions may not reach the states that need them most. If the provider network has already contracted in wage-uncompetitive states, additional funding flows to states that retained their infrastructure. Sustained rate increases, rather than temporary supplements, may be necessary to rebuild capacity in the hardest-hit markets.

Third, the concentration of effects among organizational providers highlights the importance of agency-level viability. Rate increases that flow to agencies must translate into worker compensation increases to be effective. Wage pass-through requirements—already used in some states—could ensure that Medicaid rate increases reach the front-line workforce rather than being absorbed by agency margins.

8. Conclusion

Five million Americans who depend on Medicaid home and community-based services are served by a workforce earning less than cashiers and stock clerks. This paper shows that this wage penalty is not merely unjust—it is structurally dangerous. When COVID-19 disrupted low-wage labor markets, states where home care wages were least competitive lost the most providers. The monopsony structure of Medicaid—state agencies setting rates below competitive equilibrium, with no countervailing bargaining power—created a workforce poised to collapse under stress.

The pandemic did not create the home care workforce crisis. It revealed it. The pre-existing wage gap between Medicaid-funded caregiving and available alternatives was the kindling; COVID was the match. States that maintained more competitive wages—whether through higher Medicaid rates, lower-cost-of-living labor markets, or both—entered the pandemic with a more resilient provider network and emerged with more of that network intact.

These findings point to a structural reform agenda. If Medicaid HCBS is to serve as the nation's primary system of long-term care, its reimbursement structure must account for the labor markets in which it operates. Benchmarking home care rates to local competing-sector wages would build resilience against future shocks—whether the next pandemic, the next warehouse hiring boom, or the next minimum wage increase. The fragility is a policy choice. It can be unchosen.

Acknowledgements

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

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

Contributors: Anonymous

First Contributor: https://github.com/Anonymous

References

Antonova, Sonya, Courtney Harold Van Houtven, and Samuel Shen, "Essential but Unprotected: Home Care Workers During COVID-19," *Health Affairs*, 2023, *42* (6), 823–832.

Azar, José, Ioana Marinescu, and Marshall Steinbaum, "Labor Market Concentration," *Journal of Human Resources*, 2022, *57* (S), S167–S199.

Barrero, Jose Maria, Nicholas Bloom, and Steven J Davis, "Why Working from Home Will Stick," *National Bureau of Economic Research Working Paper*, 2021, (28731).

Berger, David W, Kyle F Herkenhoff, and Simon Mongey, "Labor Market Power," *American Economic Review*, 2022, *112* (4), 1147–1193.

Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan, "How Much Should We Trust Differences-in-Differences Estimates?," *Quarterly Journal of Economics*, 2004, *119* (1), 249–275.

Bureau of Labor Statistics, "Quarterly Census of Employment and Wages," Technical Report, U.S. Department of Labor 2019. Annual average data, private sector.

Centers for Medicare and Medicaid Services, "National Plan and Provider Enumeration System (NPPES)," Technical Report, U.S. Department of Health and Human Services 2023.

_ , "T-MSIS Analytic Files (TAF)," Technical Report, U.S. Department of Health and Human Services 2024. Pre-processed analytic extract covering January 2018 through November 2024.

Chetty, Raj, John N Friedman, Nathaniel Hendren, and Michael Stepner, "How Did COVID-19 and Stabilization Policies Affect Spending and Employment? A New Real-Time Economic Tracker Based on Private Sector Data," *National Bureau of Economic Research Working Paper*, 2020, (27431).

Clemens, Jeffrey and Joshua D Gottlieb, "Do Physicians' Financial Incentives Affect Medical Treatment and Patient Health?," *American Economic Review*, 2014, *104* (4), 1320–1349.

Coibion, Olivier, Yuriy Gorodnichenko, and Michael Weber, "Labor Markets During the COVID-19 Crisis: A Preliminary View," *European Economic Review*, 2020, *128*, 103521.

Fisher, Ronald A, "The Design of Experiments," *Oliver and Boyd, Edinburgh*, 1935.

Grabowski, David C, Joseph J Angelelli, and Vincent Mor, "Medicaid Payment and Risk-Adjusted Nursing Home Quality Measures," *Health Affairs*, 2004, *23* (5), 243–252.

Howes, Candace, "Living Wages and the Retention of Home Care Workers in San Francisco," *Industrial Relations*, 2005, *44* (1), 139–163.

MACPAC, "An Updated Look at Rates of Medicaid Spending on Home and Community-Based Services," Technical Report, Medicaid and CHIP Payment and Access Commission 2022.

Manning, Alan, "Monopsony in Motion: Imperfect Competition in Labor Markets," *Princeton University Press*, 2003.

Meer, Jonathan and Jeremy West, "Effects of the Minimum Wage on Employment Dynamics," *Journal of Human Resources*, 2016, *51* (2), 500–522.

Moulton, Brent R, "An Illustration of a Pitfall in Estimating the Effects of Aggregate Variables on Micro Units," *Review of Economics and Statistics*, 1990, *72* (2), 334–338.

Ruffini, Krista, "Worker Earnings, Service Quality, and Firm Profitability: Evidence from Nursing Homes and Minimum Wage Reforms," *Review of Economics and Statistics*, 2022, pp. 1–45.

Staiger, Douglas O, Joanne Spetz, and Ciaran S Phibbs, "Do Hospitals Really Respond to Prices? The Case of Nurse Wages," *Journal of Political Economy*, 2010, *118* (6), 1169–1206.

The New York Times, "Coronavirus (COVID-19) Data in the United States," Technical Report, The New York Times 2020. GitHub repository.

U.S. Congress, "American Rescue Plan Act of 2021," Technical Report Public Law 117-2, 117th Congress 2021.

Wolfers, Justin, "Did Unilateral Divorce Laws Raise Divorce Rates? A Reconciliation and New Results," *American Economic Review*, 2006, *96* (5), 1802–1820.

Young, Alwyn, "Channeling Fisher: Randomization Tests and the Statistical Insignificance of Seemingly Significant Experimental Results," *Quarterly Journal of Economics*, 2019, *134* (2), 557–598.

A. Data Appendix

A.1 T-MSIS Data Processing

The Transformed Medicaid Statistical Information System (T-MSIS) data used in this paper are derived from a pre-processed analytic extract stored as a single Apache Parquet file (approximately 4.1 GB). The extract contains provider-month level records with the following fields: billing provider NPI, HCPCS procedure code, claim-from month, total paid amount, total claims, and total unique beneficiaries.

HCBS code selection. I define HCBS activity using HCPCS codes with T- or S-prefixes. T-codes are temporary national codes assigned by CMS specifically for Medicaid services; S-codes are temporary national codes used by Blue Cross/Blue Shield Association and adopted by state Medicaid programs. The relevant code families include:

- T1019–T1020: Personal care services
- T2016–T2017: Attendant care services
- T2022, T2025: Day habilitation, supported employment
- T2033: Residential habilitation
- S5125, S5150, S5170: Home health aide, homemaker, respite care

NPI-to-state assignment. Billing NPIs are matched to practice states using the NPPES monthly extract. For NPIs appearing in multiple states, I retain the first-listed practice address. After merging, I drop records where the NPPES state is missing or not a valid US state/DC abbreviation.

Panel construction. I aggregate to state × month by counting unique billing NPIs (providers), summing paid amounts (spending), summing claims, and summing unique beneficiaries. The panel runs from January 2018 (the earliest reliable T-MSIS month in the extract) through November 2024.

A.2 BLS QCEW Wage Data

Wage data are drawn from the Bureau of Labor Statistics Quarterly Census of Employment and Wages (QCEW) for 2019, annual averages, private-sector establishments. I access the data through the BLS QCEW public API and extract state-level average weekly wages for four NAICS industry codes:

- 624120: Services for the Elderly and Persons with Disabilities

- 445110: Supermarkets and Other Grocery Stores
- 722513: Limited-Service Restaurants
- 493110: General Warehousing and Storage

Weekly wages are converted to implied hourly rates by dividing by 40 (assuming a standard work week). The QCEW captures wages for all workers covered by unemployment insurance, representing approximately 97% of total nonfarm wage and salary civilian employment.

A.3 COVID-19 and Control Data

COVID-19 case data are from the New York Times COVID-19 tracking project GitHub repository, which compiled state-reported cumulative case counts from January 2020 through March 2023. I compute monthly new cases as the difference in cumulative counts between the last day of consecutive months, then normalize by state population from the American Community Survey. COVID case rates are set to zero for months before January 2020 (pre-pandemic) and after the end of NYT coverage (March 2023), as COVID is no longer a significant labor market factor by that date.

State unemployment rates are from the Federal Reserve Economic Data (FRED) system, Bureau of Labor Statistics Local Area Unemployment Statistics (LAUS) program, monthly, seasonally adjusted.

State population estimates (2018–2023) are from the Census Bureau American Community Survey 5-year estimates (variable B01003_001E). Population for 2024 is forward-filled from the 2023 estimates.

A.4 Sample Construction

Table 7: Sample Construction

Step	State-months	States
Raw T-MSIS HCBS panel (T/S codes)	4,233	51
After NPPES state merge	4,233	51
After wage ratio merge (drop unmatched)	4,233	51
With controls (COVID cases set to 0 pre-2020)	4,233	51

Notes: COVID case rates are set to zero for the pre-pandemic period (January 2018–December 2019) and for months after March 2023 when NYT tracking ended. Population for 2024 is forward-filled from 2023 ACS estimates. All 51 state-equivalents are retained throughout.

B. Additional Results

B.1 Raw Trends in HCBS Providers

Figure 8 shows the raw time series of log active HCBS providers for the three wage ratio terciles. All three groups show a decline beginning in early 2020, but the low-competitiveness tercile (dashed line) experiences a steeper and more sustained decline, consistent with the regression results.

Figure 8: Raw Trends in Log HCBS Providers by Wage Ratio Tercile

Notes: Average log(active HCBS providers + 1) by month for states in each tercile of the 2019 wage competitiveness ratio distribution. Vertical dashed line marks March 2020 (COVID onset).

B.2 Cross-Sectional Relationship

Figure 9 displays the cross-sectional relationship between the 2019 wage ratio and the change in log providers from the pre-COVID period (2018–2019 average) to the post-COVID period (2020–2024 average), net of state-specific trends. The positive slope confirms that wage-competitive states experienced relatively better provider supply outcomes.

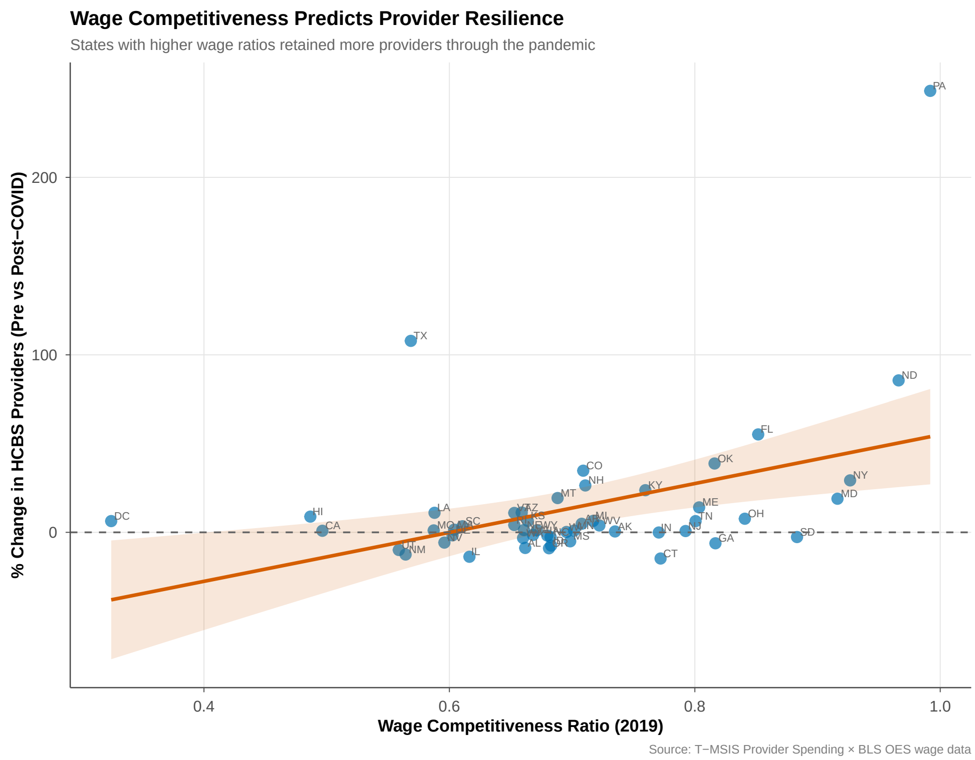

Figure 9: Wage Competitiveness and Post-COVID Provider Change

Notes: Each point represents a state. The x-axis shows the 2019 wage competitiveness ratio; the y-axis shows the change in log HCBS providers from 2018–2019 average to 2020–2024 average. The fitted line is from a bivariate OLS regression. DC is labeled as an outlier.

C. Sensitivity and Diagnostics

C.1 Tercile Event Study

Figure 6 (reproduced in the main text) shows event study coefficients from a specification interacting tercile indicators with month dummies, using the high-competitiveness tercile as the reference group. The low and medium terciles show increasingly negative coefficients after March 2020, indicating that less competitive states lost more providers relative to the most competitive states.

C.2 Behavioral Health Placebo Event Study

Figure 7 (reproduced in the main text) compares the event study coefficients for HCBS providers (T/S-codes) with those for behavioral health providers (H-codes). The divergent post-COVID patterns—HCBS showing a clear wage-ratio gradient, BH showing no gradient—support the identification assumption that the wage ratio operates through in-person care labor markets rather than generic Medicaid program effects.

C.3 Influence of Individual States

The leave-one-out analysis (Figure 4) reveals that Pennsylvania is the most influential state: excluding it reduces the coefficient to 0.436. This is consistent with Pennsylvania's large Medicaid HCBS program and its position near the median of the wage ratio distribution. Importantly, the coefficient remains positive and substantively meaningful even when Pennsylvania is excluded. The District of Columbia, despite its outlier wage ratio (0.324), does not disproportionately influence the results: excluding DC yields a coefficient of 0.795, close to the full-sample estimate.

C.4 Functional Form

The tercile specification (Table 5, Column 4) provides a non-parametric check on the linear functional form assumption. The medium-competitiveness tercile shows a significant negative coefficient relative to the high-competitiveness reference (-0.162, $p < 0.10$), while the low-competitiveness tercile shows a negative but insignificant coefficient (-0.150). The non-monotonic pattern (medium tercile more affected than low tercile) may reflect composition effects: the lowest-ratio states include large states like California and Texas where the absolute volume of HCBS activity provides some resilience.

The wage level specification (Column 5) confirms that using the PCA wage level directly (rather than the ratio) produces qualitatively similar results: a $1 increase in hourly home care wages is associated with 0.045 log points more providers ($p < 0.10$). This provides reassurance that the results are not driven by the specific construction of the ratio denominator.