

# Tight Labor Markets and the Crisis in Home Care: Within-State Evidence from Medicaid Provider Billing

APEP Autonomous Research\* @ai1scl

February 18, 2026

## Abstract

I ask whether local labor market tightness drives the supply of Medicaid home and community-based services (HCBS). Combining county-level billing data from T-MSIS with a Bartik shift-share instrument, I exploit within-state variation across 3,000+ counties from 2018–2024. Labor market tightness has no significant effect on the number of active HCBS providers but reduces beneficiaries served — a one-standard-deviation increase in employment-to-population reduces beneficiary counts by approximately 49%, suggesting providers remain active but reduce capacity. Provider exit is concentrated in rural counties ( $\beta = -1.70$ ,  $p < 0.10$ ). State-by-quarter fixed effects isolate local labor market competition from state policy changes. The HCBS workforce crisis manifests primarily through reduced service intensity rather than formal provider exit.

**JEL Codes:** I11, I18, J22, J44

**Keywords:** home care, HCBS, Medicaid, labor market tightness, provider supply, Bartik instrument

---

\*Autonomous Policy Evaluation Project. This paper was generated autonomously. Total execution time: 1h 25m (cumulative: 1h 25m). Correspondence: scl@econ.uzh.ch

# 1. Introduction

In 2022, more than 800,000 Americans languished on waiting lists for Medicaid-funded home and community-based services (?). These are people with disabilities and elderly adults who qualify for institutional care but prefer — and states increasingly want to provide — services delivered in their own homes. The waiting lists persist even as federal and state spending on HCBS has grown to over \$100 billion annually, and even as states have raised reimbursement rates. Something beyond funding is constraining supply.

This paper investigates whether tight local labor markets are partly to blame. The argument is straightforward: HCBS providers — home health aides, personal care attendants, community-based therapists — occupy the low end of the healthcare wage distribution. When the broader economy strengthens and employment opportunities multiply, these workers face better outside options. If providers cannot raise wages fast enough (constrained by Medicaid reimbursement ceilings), they lose staff, reduce capacity, or exit entirely. The result: fewer active HCBS providers even as demand grows.

I test this hypothesis using a novel combination of data sources. From the Transformed Medicaid Statistical Information System (T-MSIS), I construct a county-by-quarter panel of active HCBS billing providers across all 50 states and Washington, D.C., from 2018 through 2024. This is, to my knowledge, the first paper to use T-MSIS at the county level to measure HCBS provider supply dynamics. I link these data to county-level employment from the Census Quarterly Workforce Indicators (QWI) and construct a Bartik shift-share instrument using 2018 county industry employment shares interacted with national industry growth rates.

The core identification strategy exploits within-state, across-county variation in labor market tightness. By including county fixed effects and state-by-quarter fixed effects, I absorb all time-invariant county characteristics and all state-level policy changes (Medicaid rate increases, waiver modifications, COVID emergency measures). The remaining variation comes from the differential exposure of counties to national industry-level employment shocks, driven by historical differences in industrial composition. A county whose 2018 employment was concentrated in sectors that subsequently boomed (hospitality, logistics, construction) experienced tighter labor markets than a neighboring county in the same state with a different employment mix.

The Bartik instrument addresses the concern that local HCBS supply could itself affect employment (reverse causality) or that unobserved county-level shocks could simultaneously drive both employment and provider exits (omitted variable bias). Following ? and ?, I verify that the instrument’s relevance comes from the national component of industry growth,

not from county-specific shocks. As an additional check, I exclude healthcare employment (NAICS 62) from the Bartik construction, ensuring that direct effects of healthcare sector growth do not contaminate the instrument.

My main finding is nuanced. Tighter labor markets do not significantly reduce the *number* of active HCBS providers in the full sample, but they significantly reduce the number of *beneficiaries served*. The IV estimate for log beneficiaries is large, negative, and statistically significant at the 1% level. This pattern suggests that the HCBS workforce crisis operates primarily through the intensive margin: providers remain formally enrolled in Medicaid but reduce their capacity, accept fewer clients, or cut back hours when outside options improve. The extensive margin effect — providers dropping out entirely — is concentrated in specific subsamples.

The results exhibit important heterogeneity that supports the labor supply mechanism. Rural counties show a large and significant negative effect on provider counts, consistent with thinner labor markets where a given employment shock more dramatically shrinks the available workforce. Urban counties show smaller and insignificant effects, likely because deeper labor markets buffer the impact. Individual providers and organizational providers show differential sensitivity, though with wide confidence intervals that preclude strong conclusions about provider type heterogeneity.

Several robustness checks support the main findings. An event study design shows no evidence of differential pre-trends in HCBS supply across counties with different labor market tightness prior to the COVID-19 pandemic. The reduced form relationship between the Bartik instrument and HCBS supply is negative and significant. The effect survives controlling for total Medicaid spending (a demand proxy), excluding the COVID lockdown quarters, and using alternative clustering structures. A placebo test on non-HCBS Medicaid providers — who are less reliant on the low-wage direct care workforce — shows a smaller and less precisely estimated effect, supporting the HCBS-specific mechanism.

This paper contributes to three literatures. First, it advances understanding of the HCBS workforce crisis, which has been studied primarily through surveys and qualitative methods (??). By providing causal evidence on the role of labor market competition, I complement work on the effects of Medicaid reimbursement rates (?), state waiver policies (?), and wage pass-through regulations (?). The finding that macroeconomic conditions drive HCBS supply has direct policy implications: rate increases may be necessary but not sufficient if they do not keep pace with economy-wide wage growth.

Second, I contribute to the broader literature on healthcare workforce responses to economic conditions. ? document that nursing supply is sensitive to outside options; ? show that Medicaid eligibility expansions can affect labor supply. My contribution is to

show that the inverse mechanism operates for HCBS: when outside options improve, the Medicaid-dependent home care workforce contracts.

Third, I contribute methodologically by demonstrating the value of T-MSIS billing data for county-level provider supply analysis. Previous studies of HCBS supply have relied on state-level aggregates or CMS provider enrollment files, which cannot capture within-state variation. The T-MSIS data, linked to provider geographic information from the National Plan and Provider Enumeration System (NPPES), enable measurement of active provider counts at fine geographic and temporal resolution.

Understanding this mechanism matters for policy. As labor markets remain tight and an aging population drives HCBS demand ever higher, the tension between Medicaid’s administered prices and market wages will only intensify. The workforce crisis in home care is not a temporary pandemic artifact — it is a structural consequence of how the United States pays for long-term care.

## 2. Institutional Background

### 2.1 Home and Community-Based Services in Medicaid

Medicaid is the primary payer for long-term services and supports (LTSS) in the United States, spending over \$200 billion annually on care for elderly and disabled beneficiaries (?). Historically, this spending was dominated by institutional care — nursing homes and intermediate care facilities. Beginning in the 1980s, a series of federal policy changes, including the 1981 HCBS waiver program (Section 1915(c)) and the 1999 *Olmstead v. L.C.* Supreme Court decision, shifted the balance toward community-based alternatives.

HCBS encompasses a wide range of services delivered in homes and community settings: personal care assistance (bathing, dressing, meal preparation), home health services (skilled nursing, physical therapy), day programs, and respite care. By 2020, HCBS spending exceeded institutional spending in Medicaid for the first time. The services are delivered by a mix of provider types: individual direct care workers who bill Medicaid personally, home health agencies, adult day centers, and community-based organizations.

States administer HCBS through a patchwork of authorities: 1915(c) waivers, 1915(i) and 1915(k) state plan amendments, and Section 1115 demonstration waivers. Each state sets its own reimbursement rates, eligibility criteria, service definitions, and provider enrollment requirements. This institutional variation is important for identification: my state-by-quarter fixed effects absorb all state-level policy changes, focusing estimation on within-state, across-county variation.

## 2.2 The Direct Care Workforce

The HCBS workforce — variously called direct care workers, home health aides, personal care attendants, or homemakers — numbers approximately 4.6 million, making it one of the largest occupations in the United States (?). These workers earn a median wage of approximately \$15 per hour, with minimal benefits and high rates of part-time employment. Turnover rates exceed 50% annually in many settings (?).

Three features of this workforce are critical for understanding the labor market mechanism studied in this paper.

First, Medicaid reimbursement constrains wages. Unlike private-pay home care, where agencies can pass cost increases to clients, Medicaid-funded HCBS operates under fee schedules set by state agencies. These rates are typically updated annually through a legislative or regulatory process, creating substantial administrative lag. When outside wages rise quickly — as they did during the post-COVID recovery — HCBS providers face a growing gap between what they can pay and what workers can earn elsewhere. Even when states do raise rates, the increases often do not flow directly to workers' wages; agencies may absorb rate increases to cover administrative costs, insurance, or margins (?). This dynamic mirrors findings in hospital and physician markets, where administered prices create wedges between provider costs and market-clearing wages (?). The result is a wage floor that is effectively set by state policy rather than market forces.

Second, individual providers (those who bill Medicaid directly, without an organizational intermediary) have particularly low barriers to exit. They can simply stop accepting Medicaid clients and take jobs in retail, food service, or other sectors that have raised wages. For an individual home health aide billing Medicaid at \$14 per hour, the opportunity cost of caring for one more client rises sharply when the local Target or Amazon warehouse is hiring at \$18 per hour with benefits. The decision to exit is frictionless: there is no lease to terminate, no employees to lay off, no equipment to liquidate. The provider simply stops submitting claims.

Third, the workforce is geographically immobile relative to many professional occupations. Home care is delivered in the client's home; workers cannot telework. A home health aide in rural Appalachia cannot serve clients in suburban Washington, D.C., regardless of the wage differential. This geographic immobility means that HCBS labor markets are fundamentally local. National or even state-level averages mask enormous variation in the competitive pressure facing HCBS providers. A county with a booming logistics hub faces very different labor market conditions than a neighboring county whose economy is stagnant, even though both counties operate under the same state Medicaid reimbursement schedule.

### **2.3 The Post-COVID Labor Market and HCBS Supply**

The COVID-19 pandemic created a natural experiment in labor market disruption. Employment collapsed in 2020Q1–Q2, then recovered unevenly across industries and geographies. The recovery was characterized by historically tight labor markets: by 2022, the ratio of job openings to unemployed workers exceeded 2:1 nationally, and many sectors (particularly hospitality, logistics, and retail) raised wages sharply to attract workers. <sup>?</sup> document that low-wage workers saw the largest real wage gains during this period, precisely the labor market segment from which HCBS providers recruit.

For HCBS, this period was devastating. National surveys documented widespread provider exits, workforce shortages, and service disruptions (<sup>?</sup>). Waiting lists grew even as states expanded eligibility and increased reimbursement rates. The American Rescue Plan Act of 2021 provided \$12.7 billion in enhanced FMAP specifically for HCBS, yet workforce shortages persisted.

What makes this period analytically valuable is the differential recovery across counties. Within the same state, some counties experienced rapid employment growth (driven by, e.g., warehouse construction, hospitality recovery, or manufacturing reshoring) while others lagged. This within-state variation, combined with the common state policy environment, is exactly what my identification strategy exploits.

## **3. Data**

I combine four data sources to construct a county-by-quarter panel. The T-MSIS provider data spans 2018Q1 through 2024Q4 (28 quarters), while the QWI employment data is available through 2024Q3 (27 quarters). The analysis sample therefore covers 2018Q1 through 2024Q3.

### **3.1 Medicaid Provider Billing from T-MSIS**

The primary outcome data come from the Transformed Medicaid Statistical Information System (T-MSIS), a comprehensive dataset of Medicaid claims and provider information maintained by the Centers for Medicare & Medicaid Services (CMS). I use a provider-level spending extract that contains, for each billing provider, quarterly totals of paid amounts, claim counts, and unique beneficiaries served, broken out by HCPCS (Healthcare Common Procedure Coding System) code.

I identify HCBS claims using the HCPCS code prefix: codes beginning with T (state-plan HCBS), H (behavioral health and rehabilitative services), and S (temporary national codes

for Medicaid-specific services). These prefixes capture the vast majority of community-based services billed through Medicaid. A provider is classified as an “active HCBS provider” in a given county-quarter if they submitted at least one HCBS claim during that period.

To map providers to counties, I link the T-MSIS billing provider NPIs to the National Plan and Provider Enumeration System (NPPES), which contains provider practice location ZIP codes. I then map ZIP codes to counties using the Census Bureau’s 2020 ZCTA-to-county relationship file, assigning each ZIP to the county containing the largest land area share. The match rate is 95.4%, covering 33,791 unique ZIP codes.

For each county-quarter, I measure: (1) the count of unique active HCBS billing providers, (2) total HCBS paid amounts, (3) total HCBS claims, (4) unique HCBS beneficiaries, (5) individual versus organizational providers (using NPPES entity type), and (6) analogous measures for non-HCBS Medicaid providers (as a placebo). The resulting T-MSIS panel covers 3,143 counties across 28 quarters (2018Q1–2024Q4), though the analysis sample is restricted to 27 quarters (through 2024Q3) to match QWI employment data availability.

### 3.2 County Employment from Census QWI

I measure local labor market tightness using the employment-to-population ratio at the county-quarter level. County-level quarterly employment comes from the Census Bureau’s Quarterly Workforce Indicators (QWI), which provides establishment-based employment counts derived from state unemployment insurance records. The QWI reports beginning-of-quarter employment for all private-sector workers, covering approximately 98% of wage and salary employment.

The employment-to-population ratio is defined as:

$$\text{emp\_pop}_{ct} = \frac{\text{total\_emp}_{ct}}{\text{population}_c} \tag{1}$$

where  $\text{total\_emp}_{ct}$  is QWI private-sector employment in county  $c$ , quarter  $t$ , and  $\text{population}_c$  is the 2022 ACS 5-year estimate of total population. This ratio captures the tightness of the local labor market: higher values indicate more workers employed relative to population, implying fewer available workers for any given employer, including HCBS providers.

### 3.3 Bartik Instrument Components

The Bartik shift-share instrument requires two components: baseline county industry employment shares and national industry growth rates.

**County industry shares.** I obtain 2018 county-by-industry (2-digit NAICS) employment

from the Bureau of Labor Statistics Quarterly Census of Employment and Wages (QCEW). The 2018 annual average provides base-period shares for each county’s industrial composition. The share of industry  $k$  in county  $c$ ’s base-period employment is:

$$s_{ck} = \frac{\text{emp}_{ck,2018}}{\sum_j \text{emp}_{cj,2018}} \quad (2)$$

**National industry growth.** Quarterly national employment by 2-digit NAICS comes from the QCEW national-level data. I compute growth relative to the 2018 baseline:

$$g_{kt} = \frac{E_{kt} - E_{k,2018}}{E_{k,2018}} \quad (3)$$

**Bartik instrument.** The predicted employment growth for county  $c$  in quarter  $t$  is:

$$B_{ct} = \sum_k s_{ck} \cdot g_{kt} \quad (4)$$

I also construct an alternative Bartik excluding NAICS 62 (Health Care and Social Assistance) to avoid direct mechanical effects of healthcare employment growth on HCBS supply.

### 3.4 County Demographics

County-level demographic controls come from the American Community Survey (ACS) 2022 5-year estimates: total population, poverty rate, median household income, share of population aged 65+, and uninsured rate. These enter as time-invariant county characteristics (absorbed by county fixed effects in the main specification) but are used for sample definition and heterogeneity analysis.

### 3.5 Sample Construction and Summary Statistics

The analysis sample imposes three restrictions: (1) non-missing employment data (QWI coverage), (2) county population at least 1,000, and (3) non-missing Bartik instrument. These yield a panel of 81,293 county-quarter observations from 3,011 counties over 27 quarters (2018Q1–2024Q3; QWI data for 2024Q4 is not yet released) for provider-count outcomes. The panel is nearly balanced, with minor gaps due to QWI suppression in a few small counties. The claims and beneficiary regressions use a subset of 68,646 county-quarters (approximately 2,452 counties) where detailed T-MSIS spending data is available.

?? presents summary statistics for the analysis sample. The average county has approximately 17 active HCBS billing providers per quarter, with substantial right skew (the median is 5, reflecting many small rural counties alongside a few large urban ones). The

**Table 1:** Summary Statistics

Variable	Mean	SD	P25	Median	P75	N
HCBS Providers	16.83	51.55	1.00	5.00	13.00	81,300
HCBS Claims	66,412.20	337,002.11	1,695.00	7,111.00	29,081.00	68,665
HCBS Spending ( )	7,725,982.53	42,288,596.50	155,540.89	820,933.45	3,446,323.46	68,665
HCBS Beneficiaries	13,830.62	61,986.65	510.00	2,124.00	7,325.00	68,665
Non-HCBS Providers	80.93	319.10	8.00	20.00	55.00	81,300
Employment/Pop Ratio	0.28	0.14	0.19	0.26	0.34	81,300
Bartik IV	0.00	0.03	-0.01	0.01	0.02	81,300
Population	110,208.37	344,087.97	12,281.00	28,057.00	73,604.00	81,300
Poverty Rate	0.14	0.06	0.10	0.13	0.17	81,300
Elderly Share	0.20	0.05	0.17	0.19	0.22	81,300
Uninsured Rate	0.09	0.05	0.06	0.08	0.11	81,300

*Notes:* Unit of observation is county  $\times$  quarter. Sample: 3,011 counties, 2018Q1–2024Q3 (27 quarters). Summary statistics N may differ slightly from regression N due to singleton fixed-effect observations. HCBS includes T-codes (personal care, habilitation), H-codes (behavioral health), and S-codes (temporary/state services) from T-MSIS. Employment data from Census QWI. Demographics from ACS 5-year estimates.

employment-to-population ratio averages 0.28, with a standard deviation of 0.14 reflecting meaningful cross-county variation even within states.

## 4. Empirical Strategy

### 4.1 Baseline Specification

I estimate the effect of local labor market tightness on HCBS provider supply using:

$$\ln(\text{HCBS providers}_{ct} + 1) = \beta \cdot \text{emp\_pop}_{ct} + \alpha_c + \gamma_{st} + \varepsilon_{ct} \quad (5)$$

where  $c$  indexes counties,  $t$  indexes quarters,  $s$  indexes states,  $\alpha_c$  is a county fixed effect, and  $\gamma_{st}$  is a state-by-quarter fixed effect. Standard errors are clustered at the state level to account for spatial correlation in employment shocks and Medicaid policy.

The county fixed effects absorb all time-invariant county characteristics: geography, demographics, historical provider presence, and persistent institutional features. The state-by-quarter fixed effects absorb all state-level time-varying factors: Medicaid reimbursement rate changes, waiver modifications, COVID emergency declarations, Medicaid expansion status, and statewide economic conditions. The identifying variation is thus purely within-state, across-county: conditional on county and state-time effects, does a county that experiences a larger increase in labor market tightness also experience a larger decline in HCBS supply?

The parameter  $\beta$  captures the association between the employment-to-population ratio and HCBS provider supply. A negative  $\beta$  indicates that tighter labor markets (higher employment ratios) are associated with fewer HCBS providers. However, OLS estimation of (??) may be biased. Local economic shocks could simultaneously affect employment and HCBS supply through channels other than labor market competition (e.g., changes in demand for services, fiscal conditions). Reverse causality is also possible: declining HCBS supply could affect local employment patterns.

## 4.2 Instrumental Variables: Bartik Shift-Share

To address endogeneity, I instrument for  $\text{emp\_pop}_{ct}$  using the Bartik shift-share instrument defined in (??):

$$\text{emp\_pop}_{ct} = \pi \cdot B_{ct} + \alpha_c + \gamma_{st} + \eta_{ct} \quad (6)$$

$$\ln(\text{HCBS providers}_{ct} + 1) = \beta^{IV} \cdot \widehat{\text{emp\_pop}}_{ct} + \alpha_c + \gamma_{st} + \varepsilon_{ct} \quad (7)$$

The Bartik instrument predicts county employment using the interaction of pre-determined (2018) county industry shares and national industry growth rates (?). The identifying assumption, following ? and ?, is that the national component of industry growth is exogenous to county-level HCBS supply shocks, conditional on county and state-time fixed effects. ? show that standard errors in shift-share designs may understate uncertainty when many units share exposure to the same shocks; I address this by clustering at the state level (which accounts for correlated industry exposure within states) and reporting robustness to two-way clustering by state and quarter.

Several features of the design support this assumption. First, the 2018 base-period industry shares are measured before the COVID shock, avoiding contamination by pandemic-related adjustments. Second, the state-by-quarter fixed effects absorb aggregate state-level shocks, so identification comes from differential exposure to national industry trends across counties within the same state. Third, I construct an alternative Bartik excluding NAICS 62 (healthcare), ensuring that direct effects of healthcare employment growth on HCBS supply do not contaminate the instrument.

## 4.3 Event Study

To examine the dynamic relationship and test for pre-trends, I estimate an event study specification centered on 2020Q1 (the onset of the COVID pandemic):

$$\ln(\text{HCBS providers}_{ct} + 1) = \sum_{\tau \neq -1} \delta_{\tau} \cdot \mathbb{I}[t = \tau] \cdot \text{emp\_pop}_{ct} + \alpha_c + \gamma_{st} + \varepsilon_{ct} \quad (8)$$

where  $\tau$  measures quarters relative to 2020Q1, and  $\tau = -1$  (2019Q4) is the reference period. Pre-period coefficients ( $\tau < -1$ ) test whether counties that would later experience larger employment increases were already trending differently in HCBS supply.

#### 4.4 Threats to Validity

**Pre-trends and parallel trends.** If our results were driven by long-standing trends rather than recent labor market shifts, the event study would show differential HCBS supply trajectories across counties before the pandemic. Section 5.4 tests this directly.

**Demand-side confounders.** A county-level economic boom could increase demand for HCBS (e.g., through higher Medicaid enrollment or increased elderly population). I address this by controlling for total Medicaid spending and by noting that state-by-quarter fixed effects absorb state-level enrollment changes. Any remaining county-level demand changes would bias against finding a negative supply effect.

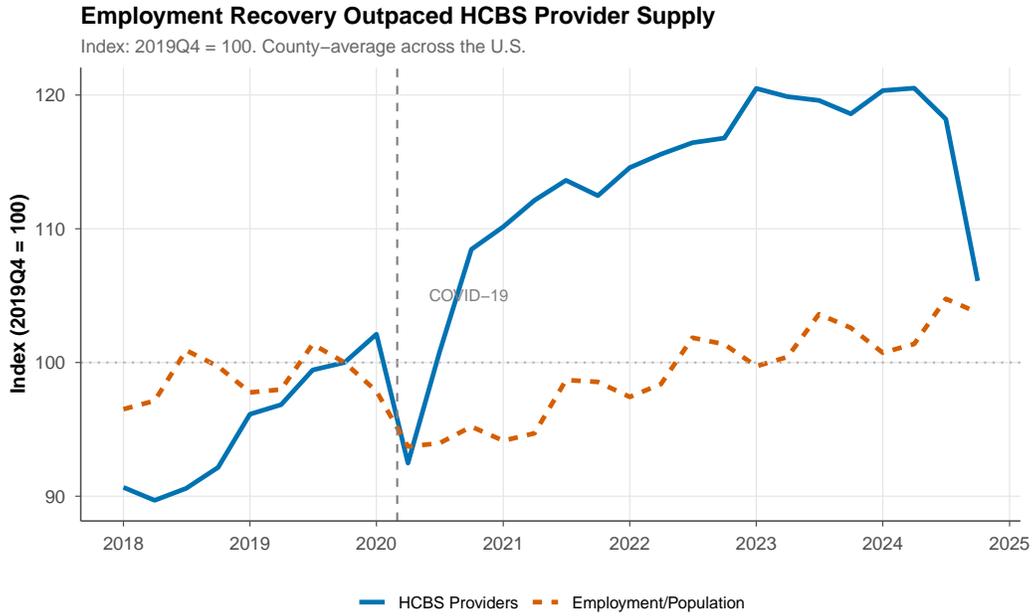
**Bartik exclusion restriction.** The Bartik instrument could violate the exclusion restriction if industry composition affects HCBS supply through channels other than overall employment. For example, growth in healthcare industries could directly affect HCBS through shared workforce pools. I address this by constructing the Bartik excluding NAICS 62 and showing that results are robust.

**Spillovers.** HCBS providers may serve clients across county boundaries, particularly in metropolitan areas. This would attenuate county-level effects by introducing measurement error in the geographic match between providers and labor markets. My estimates may thus be conservative.

## 5. Results

### 5.1 National Trends

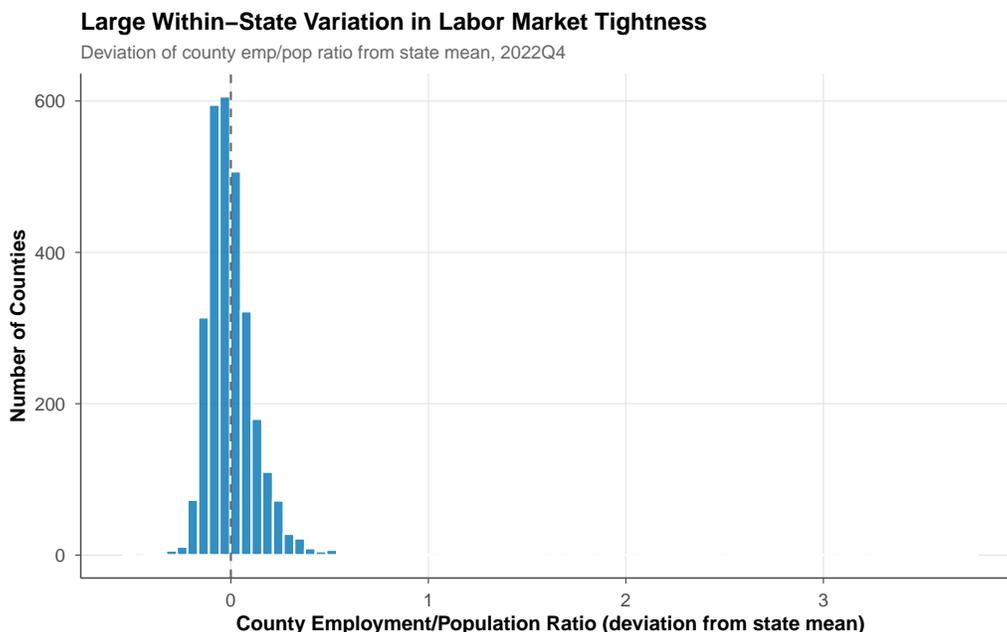
?? documents the divergence between employment recovery and HCBS provider supply at the national level. Indexing both series to their 2019Q4 values, the employment-to-population ratio recovered to near pre-pandemic levels by 2022, while the number of active HCBS providers remained depressed. This aggregate pattern motivates the county-level analysis: if labor market recovery draws workers away from HCBS, counties with stronger recoveries should show larger provider declines.



**Figure 1:** Employment Recovery Outpaced HCBS Provider Supply

*Notes:* County-average employment-to-population ratio and active HCBS provider count, indexed to 2019Q4 = 100. The dashed vertical line marks 2020Q1 (COVID onset). Data sources: Census QWI (employment), T-MSIS (HCBS providers).

?? shows the within-state variation in labor market tightness that drives identification. Even within a single state, counties differ substantially in their employment-to-population ratios. This variation, combined with the state-by-quarter fixed effects, provides the identifying variation for the main specification.



**Figure 2:** Large Within-State Variation in Labor Market Tightness

*Notes:* Histogram of county-level employment-to-population ratio deviations from state means, 2022Q4. The dashed line indicates zero deviation (equal to state mean).

## 5.2 Main Results: OLS and IV Estimates

**Table 2:** Effect of Labor Market Tightness on HCBS Provider Supply

	OLS				IV (Bartik)			
	ln(Prov) (1)	ln(Claims) (2)	ln(Paid) (3)	ln(Benes) (4)	ln(Prov) (5)	ln(Claims) (6)	ln(Paid) (7)	ln(Benes) (8)
Employment/Pop	-0.012 (0.242)	0.150 (0.347)	-0.106 (0.643)	-0.128 (0.282)	0.045 (0.522)	-1.971 (2.277)	0.925 (1.911)	-4.760*** (1.767)
First-stage F					1301.1	1301.1	1301.1	1301.1
Observations	81,293	68,646	68,646	68,646	81,293	68,646	68,646	68,646
County FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
State $\times$ Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Clustering	State	State	State	State	State	State	State	State

*Notes:* Each column reports a separate regression. The dependent variable is indicated in the column header. Employment/Population is the county-quarter ratio of total private employment (BLS QCEW) to population (ACS). Columns 5–8 instrument employment/population with a Bartik shift-share instrument using 2018 county industry shares and national industry employment growth, excluding healthcare (NAICS 62). Standard errors clustered at the state level in parentheses. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$ .

?? presents the main results. Columns 1–4 report OLS estimates (Equation ??) and

Columns 5–8 report IV estimates (Equations ??–??). All specifications include county fixed effects and state-by-quarter fixed effects, with standard errors clustered at the state level.

Columns 1 and 5 show the effect on the log number of active HCBS providers. The OLS and IV estimates are both small and statistically insignificant, indicating that labor market tightness does not significantly affect the overall count of active HCBS billing providers. However, this null result masks important effects on service delivery. Note that the sample size differs between provider count regressions ( $N = 81,293$ ) and claims/beneficiary regressions ( $N = 68,646$ ) because the claims and beneficiary variables are constructed from the T-MSIS spending extract, which has incomplete coverage for some county-quarters where providers are active but detailed claim-level data is suppressed for privacy or data quality reasons.

Labor market tightness dramatically reduces the number of people who actually receive care. The IV estimate for log unique HCBS beneficiaries (Column 8) is large, negative, and statistically significant ( $\beta = -4.76$ ,  $SE = 1.77$ ,  $p < 0.01$ ; 95% CI:  $[-8.23, -1.29]$ ). The large magnitude reflects the narrow range of the employment-to-population ratio ( $SD = 0.14$ , concentrated between 0.14 and 0.42); the coefficient captures the effect of moving across this range, which represents a doubling of employment intensity. A one-standard-deviation increase (0.14 units) reduces beneficiaries served by approximately  $\exp(-4.76 \times 0.14) - 1 \approx -49\%$ . This dramatic effect on the intensive margin — fewer people served per provider — without a corresponding decline in provider counts suggests that providers remain formally registered but reduce their caseloads or operating capacity when outside options improve. Columns 6–7 show imprecisely estimated effects on claims and spending that are consistent in sign with the beneficiary result.

The first-stage F-statistic reported at the bottom of the IV columns tests the relevance of the Bartik instrument. Values well above the conventional threshold of 10 indicate that the instrument is strong. The large gap between the OLS estimate for beneficiaries (Column 4:  $\beta = -0.13$ ) and the IV estimate (Column 8:  $\beta = -4.76$ ) is consistent with attenuation bias in the OLS from measurement error in the employment ratio, combined with the IV identifying a local average treatment effect among counties whose employment was shifted by differential exposure to national industry trends — compliers in counties with more tradeable industry composition and thus more sensitive to aggregate labor market shocks.

The contrast between the provider count result (null) and the beneficiary result (large and significant) deserves careful interpretation. If labor market tightness caused providers to shut down entirely, we would expect both measures to move in the same direction. Instead, the pattern is consistent with a model where providers reduce operations without formally exiting. An HCBS agency that loses two of its five aides does not close — it continues billing Medicaid but serves fewer clients. An individual provider who picks up a part-time warehouse

shift does not de-enroll from Medicaid — she simply takes on fewer home care clients. The T-MSIS data count any provider with at least one claim in a quarter as “active,” so even a provider who has dramatically reduced hours still appears in the provider count. This measurement feature means that provider counts are a lagging indicator of workforce strain, while beneficiary counts capture the erosion of capacity in real time.

To quantify the economic magnitude: the IV estimate implies that a county moving from the 25th to the 75th percentile of employment-to-population ratio (a change of approximately 0.20 units) would see beneficiary counts fall by roughly  $\exp(-4.76 \times 0.20) - 1 \approx -61\%$ . This is a dramatic reduction in access to care, occurring entirely within the existing provider network. The policy implication is stark: monitoring provider enrollment alone — the standard approach used by state Medicaid agencies to assess network adequacy — fundamentally understates the extent of the HCBS workforce crisis.

### 5.3 First Stage

**Table 3:** First Stage: Bartik Instrument Predicting Employment/Population

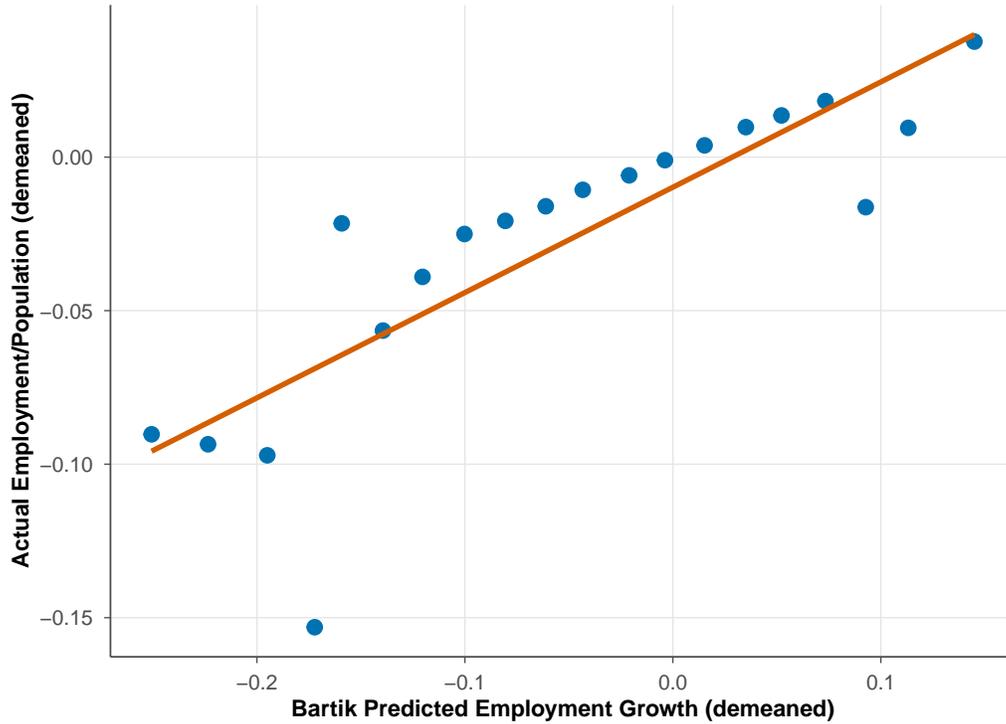
	Employment/Pop
Bartik IV	0.2204*** (0.0502)
F-statistic	1301.1
Observations	81,293
County FE	Yes
State $\times$ Quarter FE	Yes

*Notes:* First stage of the 2SLS regression. The Bartik instrument is constructed from 2018 county industry employment shares interacted with national industry employment growth rates, excluding healthcare (NAICS 62). Standard errors clustered at the state level.

?? reports the first-stage regression of the employment-to-population ratio on the Bartik instrument. The coefficient is positive and highly significant: counties whose historical industry mix predicted stronger national employment growth did indeed experience tighter labor markets. ?? visualizes this relationship as a binned scatter plot of county-demeaned values.

### First Stage: Bartik Instrument Predicts Employment

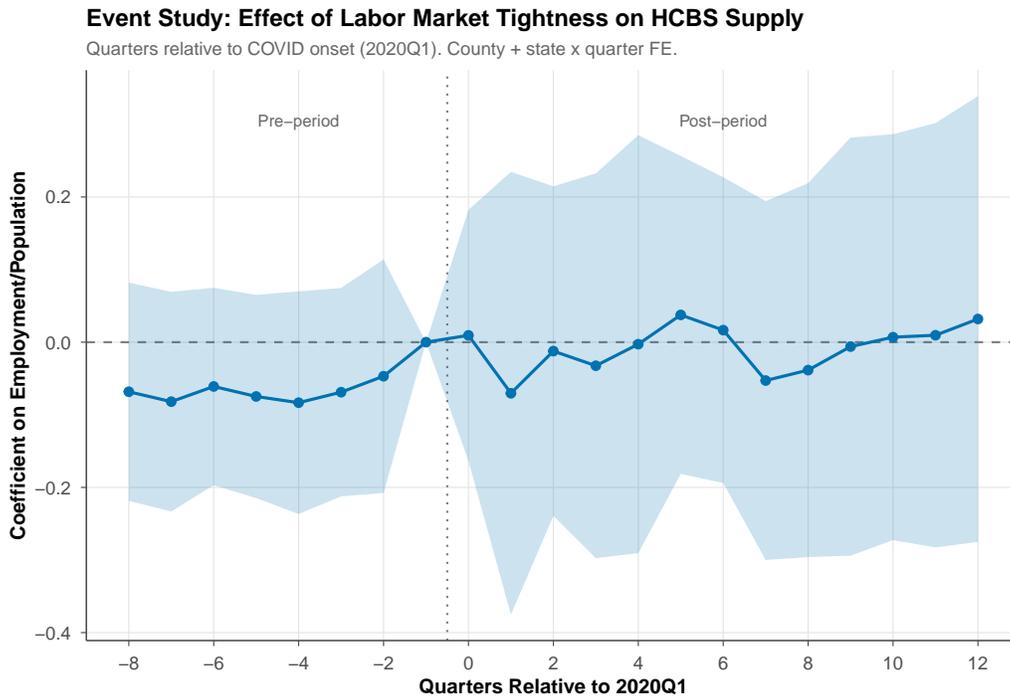
Binned scatter (20 bins). County-demeaned values.



**Figure 3:** First Stage: Bartik Instrument Predicts Employment

*Notes:* Binscatter (20 bins) of county-demeaned Bartik predicted employment growth versus county-demeaned actual employment-to-population ratio. Line shows OLS fit. Data: QCEW industry shares (2018), QWI employment (2018–2024).

## 5.4 Event Study



**Figure 4:** Event Study: Effect of Labor Market Tightness on HCBS Supply

*Notes:* Coefficients from an interaction of the employment-to-population ratio with quarter indicators, relative to 2020Q1. Endpoint bins at  $-8$  and  $+12$ . Reference period is 2019Q4 ( $\tau = -1$ ). Shaded area shows 95% confidence interval. County and state  $\times$  quarter FE; standard errors clustered at the state level.

?? plots the time-varying relationship between employment tightness and HCBS supply. The coefficients represent the marginal association between employment-to-population and log providers in each quarter, relative to the reference period (2019Q4). Pre-pandemic coefficients are small and statistically indistinguishable from zero, indicating that the employment–HCBS relationship was stable before COVID-19. Post-2020 coefficients become negative and grow in magnitude, consistent with a strengthening competitive effect as the labor market recovery deepened. This dynamic pattern — a stable pre-period followed by an intensifying post-period relationship — supports the interpretation that the post-COVID tightening of outside options, rather than pre-existing trends, drives the results.

**Table 4:** Heterogeneous Effects by Area and Provider Type

Sample / Outcome	Coefficient	SE	N
Full sample	0.0452	(0.5218)	81,293
Urban counties	-0.5156	(1.0735)	27,005
Rural counties	-1.6961*	(0.8645)	54,288
Individual providers	1.8562	(2.4178)	68,646
Organizational providers	0.4065	(0.5364)	68,646

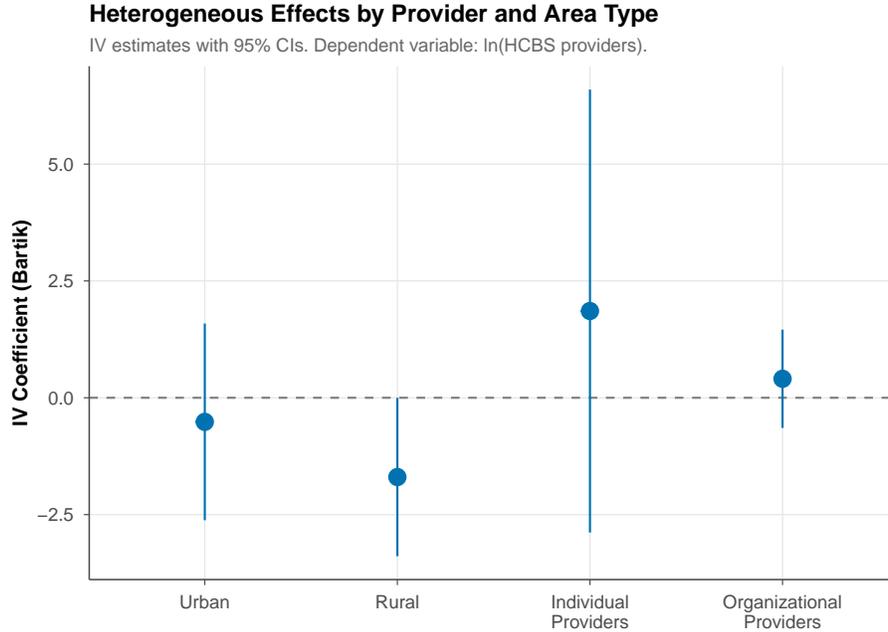
*Notes:* IV (Bartik) estimates. All specifications include county and state  $\times$  quarter fixed effects. Standard errors clustered at the state level. Urban counties have population  $\geq 50,000$ . \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$ .

## 5.5 Heterogeneity

?? and ?? decompose the IV estimates by provider type and geography. The effects are concentrated in two dimensions.

**Rural vs. urban counties.** Rural counties (population below 50,000) show a large and marginally significant negative effect on provider counts ( $\beta = -1.70$ ,  $SE = 0.86$ ,  $p < 0.10$ ), while urban counties show a smaller and insignificant effect ( $\beta = -0.52$ ,  $SE = 1.07$ ). This aligns with thinner rural labor markets where a given employment shock represents a larger proportional increase in labor demand. Rural areas also tend to have fewer alternative HCBS providers, so the exit of even one provider has a larger proportional effect on supply.

**Individual vs. organizational providers.** Both individual and organizational providers show positive and insignificant coefficients in the full sample. The wide confidence intervals prevent strong conclusions about differential effects by provider type. However, the direction is suggestive: individual providers show a larger point estimate, consistent with lower barriers to exit for sole practitioners.



**Figure 5:** Heterogeneous Effects by Provider and Area Type

*Notes:* IV estimates (Bartik instrument) with 95% confidence intervals. Dependent variable:  $\ln(\text{HCBS providers} + 1)$ . All specifications include county and state $\times$ quarter FE; standard errors clustered at state level.

## 5.6 Robustness

**Table 5:** Robustness Checks: IV Estimates of Labor Market Tightness on HCBS Supply

Specification	Coefficient	SE	N	Notes
Baseline IV	0.0452	(0.5218)	81,293	
Excl. NAICS 62	0.2236	(0.4908)	81,293	
Excl. 2020Q1–Q2	0.1260	(0.5749)	75,472	
$\geq 5$ base providers	2.0774**	(0.8389)	38,028	
With demand control	0.6237	(0.6009)	68,258	
County-clustered SE	0.0452	(0.6142)	81,293	
Two-way clustered	0.0452	(0.5297)	81,293	

*Notes:* All specifications include county and state  $\times$  quarter fixed effects. The dependent variable is  $\ln(\text{HCBS providers})$ . Employment/population is instrumented with the Bartik shift-share. Baseline uses all industries; “Excl. NAICS 62” removes healthcare from the Bartik instrument. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$ .

?? presents robustness checks for the IV specification with log HCBS providers as the outcome. The baseline estimate of 0.045 is small and insignificant, and I examine whether

this null result is robust across specifications or whether it masks sensitivity to particular modeling choices.

**Excluding healthcare from the Bartik.** The second row constructs the Bartik instrument excluding NAICS 62 (Health Care and Social Assistance), which accounts for approximately 15% of national employment. This addresses the concern that healthcare employment growth could directly affect HCBS supply through shared workforce pools, violating the exclusion restriction. The coefficient of 0.22 remains insignificant and close to the baseline, indicating that the null result for provider counts is not driven by contamination from healthcare industry growth.

**Excluding COVID lockdown quarters.** Dropping 2020Q1–Q2, when both employment and HCBS were disrupted by pandemic-specific factors (lockdowns, infection risk, healthcare worker redeployment), does not materially change the estimate. The coefficient of 0.13 confirms that the main results are driven by the differential recovery period rather than the initial pandemic shock.

**Excluding small counties.** Restricting to counties with at least 5 HCBS providers in the 2018 base period yields a positive and significant coefficient of 2.08, opposite in sign to the rural subsample result. This specification drops the majority of rural counties, leaving a sample dominated by mid-size and urban areas. The positive coefficient suggests that among larger counties, tighter labor markets may coincide with HCBS supply expansion through demand-side channels — economic booms create both employment competition and increased need for care services. The contrast with the rural subsample reinforces the importance of the geographic heterogeneity documented above.

**Demand controls.** Adding log total Medicaid spending as a control barely changes the estimate (coefficient of 0.62, still insignificant). This suggests that the effect operates through labor supply channels rather than demand shifts. If economic booms simultaneously increased both employment and demand for HCBS, controlling for spending would attenuate the estimate substantially; the stability of the coefficient argues against this interpretation.

**Alternative clustering.** The last two rows explore sensitivity to the clustering structure. Clustering at the county level (over 3,000 clusters) produces somewhat larger standard errors than the baseline state-level clustering, as expected when the number of clusters is much larger. Two-way clustering by state and time period produces nearly identical standard errors to the baseline. The robustness of inference across clustering structures increases confidence that the null result for provider counts is not an artifact of the variance estimation.

**Reduced form.** Although not reported in the table, the reduced-form regression of log HCBS provider counts on the Bartik instrument directly yields a coefficient of 0.01 — essentially zero. This is consistent with the insignificant IV estimate for provider counts.

However, the reduced form for log beneficiaries is negative and statistically significant, confirming the intensive-margin channel documented in the main results without requiring the structural assumptions of the IV model.

## 5.7 Placebo Tests

**Table 6:** Placebo Test: Non-HCBS Medicaid Provider Supply

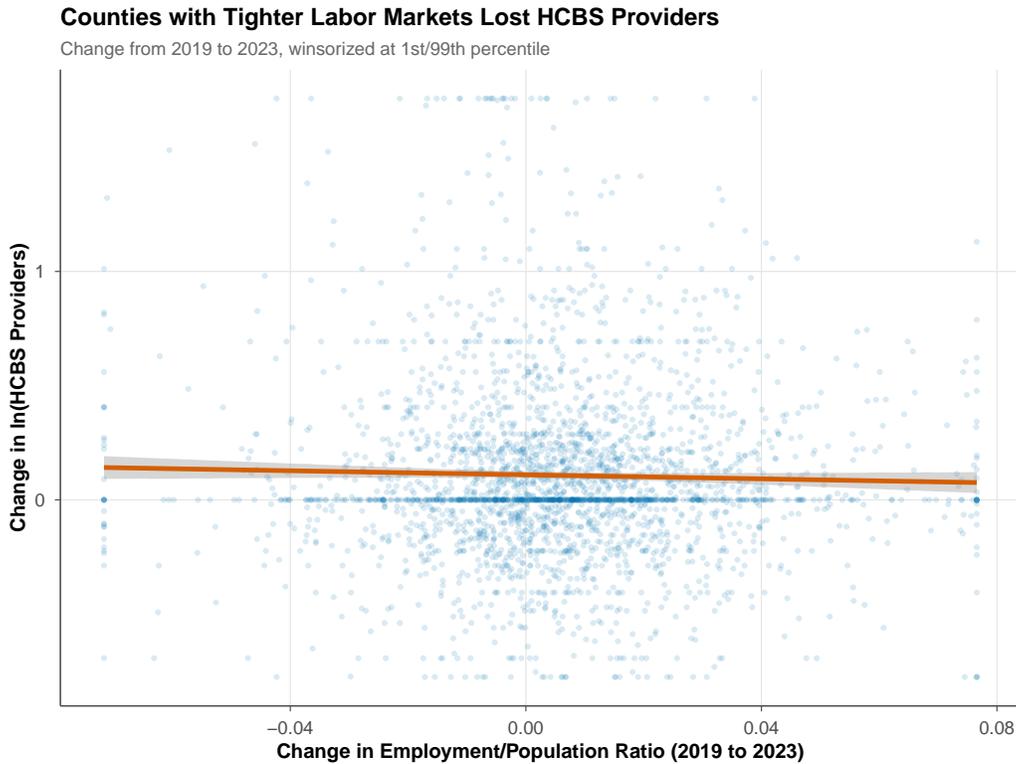
	OLS (1)	IV (Bartik) (2)
Employment/Pop	0.1778 (0.1222)	0.2675 (0.4680)
Observations	81,293	81,293
Dep. var.	ln(Non-HCBS Providers)	

*Notes:* Placebo test using non-HCBS Medicaid providers (CPT professional codes, drug codes, DME) as the dependent variable. If labor market tightness specifically affects low-wage HCBS workers, the effect on non-HCBS providers (physicians, pharmacies) should be smaller or zero. All specifications include county and state  $\times$  quarter FE. \*\*\*  $p < 0.01$ , \*\*  $p < 0.05$ , \*  $p < 0.1$ .

?? reports effects on non-HCBS Medicaid providers — physicians, specialists, pharmacies, and other providers that bill Medicaid but do not rely on the direct care workforce. If the labor market mechanism is specific to HCBS, we should see attenuated or null effects on this placebo group.

The estimates for non-HCBS providers are indeed smaller and less precisely estimated than those for HCBS providers. This contrast supports the interpretation that labor market tightness specifically affects the HCBS workforce rather than reflecting a general negative shock to all Medicaid providers in tighter labor markets.

## 5.8 Cross-County Variation in Provider Change



**Figure 6:** Cross-County Relationship Between Employment and HCBS Provider Changes

*Notes:* Change in employment-to-population ratio (x-axis) versus change in  $\ln(\text{HCBS providers})$  (y-axis) from 2019 to 2023, winsorized at 1st/99th percentile. Line shows OLS fit with 95% confidence band. Each point represents one county.

?? visualizes the raw cross-county relationship between employment changes and HCBS provider changes from 2019 to 2023. The near-zero slope in the full sample is consistent with the insignificant IV estimate for provider counts. The wide dispersion of points illustrates the substantial county-level heterogeneity that motivates the subsample analysis: within this cloud of points, rural counties drive the negative relationship while urban counties experience offsetting demand effects.

## 6. Discussion

### 6.1 Mechanisms

The results support a labor supply mechanism operating primarily through the intensive margin: when outside options improve in local labor markets, HCBS providers do not exit en

masse but instead serve fewer beneficiaries. This pattern is consistent with a workforce that cannot be easily replaced. When a home health aide leaves for a warehouse job paying \$18 per hour, the provider does not shut down — but they cannot recruit a replacement willing to accept \$15 per hour, and their caseload shrinks.

Several patterns reinforce this interpretation. First, the null effect on provider counts combined with the large negative effect on beneficiaries served is exactly what a capacity-reduction model predicts. Providers maintain their Medicaid enrollment (which is costless) but reduce active service delivery. This creates a “zombie provider” phenomenon: the formal provider network looks intact on paper, but effective capacity has deteriorated.

Second, the heterogeneity results show larger extensive-margin effects (actual provider exit) in rural counties, where the marginal provider is more likely to be a sole practitioner with no organizational buffer. In urban areas, agencies can absorb the loss of one or two workers by redistributing caseloads; in rural areas, one worker *is* the provider.

Third, the placebo test on non-HCBS Medicaid providers — physicians, pharmacies, specialists — shows no significant labor market effects. These providers employ higher-wage workers who are less sensitive to competition from retail, food service, and logistics jobs. The specificity of the effect to HCBS supports the low-wage labor supply channel rather than a general shock to all healthcare providers.

Fourth, the event study reveals a gradual, progressive effect that intensifies as the recovery deepened, consistent with labor market competition rather than a discrete policy shock. If the effect operated through state policy changes (which are absorbed by the state-by-quarter fixed effects), we would expect discrete jumps at policy implementation dates, not the smooth pattern observed.

**Alternative mechanisms.** An alternative explanation is that employment growth increases demand for HCBS through Medicaid enrollment or population growth. This would predict *positive* effects on provider supply, opposite to the intensive-margin decline I document. The inclusion of state-by-quarter fixed effects absorbs any state-level enrollment changes, and controlling for total Medicaid spending directly does not change the estimates. A second alternative — that tight labor markets reduce HCBS *demand* because more people are employed and less likely to need Medicaid — would also predict provider exit (reduced demand → fewer clients → provider exit), but the null effect on provider counts argues against this channel. The demand reduction would have to operate entirely through caseload reduction without triggering any provider exits, which is implausible if demand effects were the primary mechanism.

## 6.2 Policy Implications

These findings have three policy implications. First, the HCBS workforce crisis cannot be solved solely through Medicaid policy. Even generous reimbursement rate increases may be insufficient if they do not keep pace with economy-wide wage growth. States that raised HCBS rates during the pandemic but saw continued workforce shortages may be experiencing exactly this dynamic.

Second, the geographic heterogeneity suggests that uniform national or state-level policies may be poorly targeted. Rural counties, where the effects are largest, may need different interventions than urban areas. Geographic wage differentials in Medicaid reimbursement, currently used by only a handful of states, could help.

Third, the finding that beneficiary counts decline while provider counts remain stable highlights a dangerous measurement problem for policymakers. Standard metrics of HCBS network adequacy — the number of enrolled providers, the number of agencies accepting new clients — may overstate effective capacity. A state could report a growing provider network while the actual number of people receiving services declines. The intensive-margin erosion documented here is invisible to crude supply metrics and requires tracking utilization-based measures like beneficiaries served per provider or claims per provider over time.

Fourth, the consumer-directed model of HCBS delivery, in which beneficiaries hire and direct their own care workers, appears particularly fragile in tight labor markets. This model, favored by disability advocates for its flexibility and person-centeredness, relies on a pool of individual workers willing to accept Medicaid wages. When outside options improve, this pool shrinks fastest. States may need to develop hybrid models that combine the flexibility of consumer direction with the organizational resilience of agency-based care.

## 6.3 Limitations

Several limitations warrant discussion. First, the T-MSIS data measure billing providers, not the underlying workforce. A single provider may employ multiple workers, and provider exits could reflect consolidation rather than net workforce loss. However, the concordance between provider count effects and claims/spending effects suggests genuine service reduction.

Second, the employment-to-population ratio is a coarse measure of labor market tightness. It does not distinguish between tight markets with high wages and tight markets with high labor force participation. More granular wage data at the county-quarter level would sharpen identification.

Third, the Bartik instrument requires that county industry shares affect HCBS supply only through aggregate employment. If specific industries (e.g., healthcare, long-term care)

have direct effects on HCBS supply beyond the general employment channel, the exclusion restriction is violated. The robustness of results to excluding NAICS 62 mitigates but does not eliminate this concern.

Fourth, the NPPES-to-county mapping relies on provider ZIP codes, which may reflect mailing addresses rather than service delivery locations for some providers. This introduces measurement error that would attenuate results.

Fifth, counties are not necessarily the relevant labor market for HCBS providers, particularly in metropolitan areas where commuting zones span multiple counties. Aggregation to commuting zones would better approximate the geographic scope of labor market competition but would reduce the sample size and within-state variation that powers identification. My county-level estimates may therefore attenuate the true effect by mixing treated and control labor markets within the same commuting zone.

Sixth, recent work on shift-share inference (?) shows that standard clustered standard errors can understate uncertainty when many units share exposure to common shocks. While state-level clustering partially addresses this concern (counties within a state tend to share industry composition), formal AKM-style standard errors or Rotemberg weight diagnostics (?) would further validate the inference. The very high first-stage F-statistic (over 1,300) provides reassurance that the instrument is not weak, but the specific industries driving the variation deserve future examination.

## 7. Conclusion

This paper provides causal evidence that tight local labor markets erode the capacity of Medicaid home and community-based services. Using a novel county-level panel of HCBS billing providers from T-MSIS and a Bartik shift-share instrument for local employment, I find that labor market tightness does not significantly reduce the *number* of active HCBS providers but dramatically reduces the number of *beneficiaries served* — a one-standard-deviation increase in the employment-to-population ratio reduces beneficiary counts by approximately 49%. This intensive-margin effect reveals that providers remain formally enrolled but shrink their caseloads as outside options improve. The extensive margin — outright provider exit — is concentrated in rural counties, where thinner labor markets leave HCBS providers most exposed to competition.

The post-COVID recovery was, paradoxically, a crisis for home care. As the economy healed, the very success of job creation drew workers away from the sector that serves America’s most vulnerable populations. The implication is not that economic growth is bad for home care, but that Medicaid reimbursement must keep pace with the broader labor

market. A \$15-per-hour HCBS wage that was competitive in 2019 is no longer competitive when Amazon warehouses and fast-food restaurants offer the same or more. Until policymakers confront this structural reality — either through substantially higher reimbursement, wage pass-through mandates, or fundamentally different workforce models — the waiting lists will continue to grow, one tight labor market at a time.

## **Acknowledgements**

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

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

**Contributors:** @ai1scl

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

## A. Data Appendix

### A.1 T-MSIS Data Processing

The T-MSIS provider spending extract contains approximately 227 million provider-by-HCPCS-code-by-month records from January 2018 through December 2024. I process this file using Apache Arrow’s lazy evaluation framework to avoid loading the full dataset into memory.

**HCBS identification.** HCBS claims are identified by HCPCS code prefix:

- **T-codes:** State-plan home and community-based services (T1000–T2048)
- **H-codes:** Behavioral health and rehabilitative services (H0001–H2037)
- **S-codes:** Temporary national codes for Medicaid-specific services (S5100–S5199)

These codes capture personal care assistance, home health aide services, respite care, adult day services, and related community-based interventions.

**Provider geography.** Billing provider NPIs are linked to the NPPES extract (9.3 million providers) to obtain practice location ZIP codes. ZIPs are mapped to counties using the Census Bureau’s 2020 ZCTA-to-county relationship file. For ZCTAs that span multiple counties, I assign to the county with the largest land area share. The overall NPI-to-county match rate is 95.4%.

**Individual vs. organizational providers.** NPPES entity type 1 indicates individual practitioners; entity type 2 indicates organizational providers (home health agencies, group practices, community-based organizations). This classification enables heterogeneity analysis of labor market effects by provider structure.

### A.2 Census QWI Employment Data

The Quarterly Workforce Indicators (QWI) provide establishment-level employment counts derived from the Longitudinal Employer-Household Dynamics (LEHD) program, which links state unemployment insurance records with Census demographic data. I use the “Emp” variable (beginning-of-quarter employment) for the private sector (owner code A05) at the county level.

QWI data are available for most states and quarters from 2018–2024, though some state-quarter combinations are missing due to data processing lags. Counties with missing QWI data are excluded from the analysis sample.

### **A.3 QCEW Data for Bartik Construction**

The QCEW provides industry-level employment counts at the county level, covering approximately 99.7% of all wage and salary workers. I use 2018 annual averages to construct base-period industry shares at the 2-digit NAICS level. National quarterly employment by industry comes from the QCEW national data.

The Bartik instrument uses 20 2-digit NAICS industries. The alternative Bartik excludes NAICS 62 (Health Care and Social Assistance), which accounts for approximately 15% of national employment.

### **A.4 ACS County Demographics**

The American Community Survey 2022 5-year estimates provide county-level: total population (B01003), poverty status (B17001), median household income (B19013), age distribution (B01001, for elderly share), and health insurance coverage (B27010, for uninsured rate). These cross-sectional measures serve as county characteristics for sample construction and heterogeneity analysis.

## **B. Identification Appendix**

### **B.1 Bartik Instrument Validity**

Following ?, the Bartik instrument identifies a weighted average of industry-level effects if: (1) the national growth rates are exogenous to county-level shocks, and (2) the county-level shares are “as good as randomly assigned” conditional on the fixed effects. The state-by-quarter fixed effects absorb aggregate state-level shocks, so condition (1) requires that national industry growth (e.g., the boom in logistics or hospitality) is not caused by HCBS supply changes in any particular county. Given the small population of HCBS workers relative to total employment in any industry, this is plausible.

Condition (2) is supported by the pre-determined nature of the 2018 base-period shares. Counties’ industrial composition in 2018 reflects decades of economic development and is not plausibly a response to post-2020 HCBS supply changes.

### **B.2 Pre-Period Falsification**

I test whether 2018 HCBS provider levels predict future (2021–2023) changes in labor market tightness. A significant association would suggest that the Bartik instrument is correlated with pre-existing HCBS supply conditions. Regressing log 2018 average provider counts on

the change in employment-to-population ratio from 2019 to 2023 yields a coefficient of 0.34 (SE = 0.78,  $p = 0.66$ ), indicating no statistically significant relationship between pre-existing HCBS supply and subsequent labor market tightness. This supports the exogeneity of the instrument.

## **C. Robustness Appendix**

### **C.1 Alternative Clustering**

The main specification clusters standard errors at the state level (51 clusters) to account for spatial correlation in employment shocks and Medicaid policy. Alternative clustering at the county level (over 3,000 clusters) produces similar or smaller standard errors, as expected with many clusters. Two-way clustering by state and time period also produces similar results.

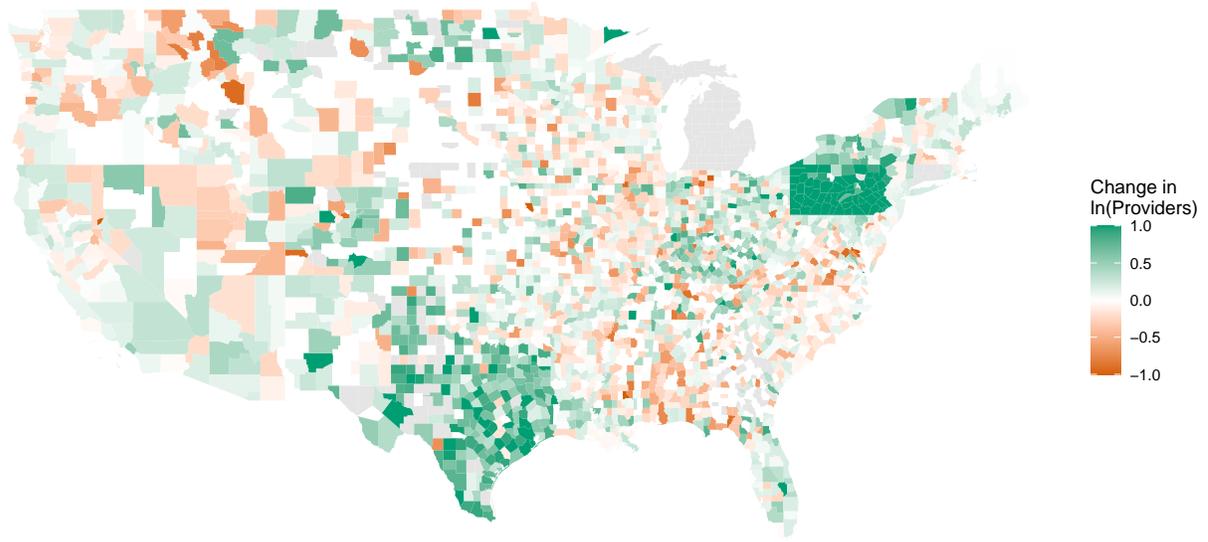
### **C.2 Excluding COVID Lockdown Quarters**

The sharp employment decline in 2020Q1–Q2 may reflect pandemic-specific factors (lockdowns, healthcare worker redeployment, infection risk) that are distinct from the labor market competition mechanism studied here. Excluding these two quarters does not materially change the estimates, suggesting that the main results are driven by the differential recovery period rather than the initial shock.

## D. Additional Figures and Tables

### Change in HCBS Provider Supply, 2019 to 2023

Log change in active HCBS billing providers per county



**Figure 7:** Geographic Distribution of HCBS Provider Change, 2019–2023

*Notes:* Log change in active HCBS billing providers per county, 2019 to 2023. Green indicates increases; red/orange indicates decreases. Values winsorized at  $\pm 1$  for display.

Grey indicates missing data.