The Depleted Safety Net: Hysteresis in Medicaid's Home Care Workforce*

APEP Autonomous Research†
@SocialCatalystLab

February 26, 2026

Abstract

Do safety-net labor markets self-correct, or does gradual erosion create fragility that only a crisis reveals? Using provider-level Medicaid claims for 617,000 billing entities across all 50 states over seven years, I show that states where more HCBS providers exited the workforce before COVID-19 experienced 6 percent larger supply declines and 7 percent larger reductions in beneficiary access after March 2020. Pre-pandemic provider depletion strongly predicts pandemic-era service disruption. Despite $37 billion in American Rescue Plan investment, the most depleted states had not recovered by December 2024. The results are consistent with hysteresis in safety-net labor markets: gradual erosion weakened provider networks, the pandemic exposed this fragility, and the system did not bounce back.

JEL Codes: I11, I13, I18, J44, H75
Keywords: Medicaid, HCBS, hysteresis, provider supply, COVID-19, safety net, workforce

*This paper is a revision of APEP-0454. See https://github.com/SocialCatalystLab/ape-papers/tree/main/apep_0454 for prior versions.

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

1. Introduction

A basic prediction of competitive labor markets is self-correction. Workers exit; wages adjust; new workers enter. The equilibrium may shift, but the market clears. This paper asks whether safety-net labor markets work that way—and finds evidence that they may not.

The setting is Medicaid's home and community-based services (HCBS) workforce: the personal care aides, behavioral health counselors, and home health workers who deliver hands-on care to approximately 5 million elderly and disabled Americans. Between 2018 and early 2020, this workforce was quietly eroding. Providers were exiting the Medicaid billing system at rates that varied dramatically across states—from 14 percent of the active provider base in some states to over 50 percent in others. Then COVID-19 arrived, and the erosion became a crisis.

Using the newly released Transformed Medicaid Statistical Information System (T-MSIS) provider spending dataset—the first public provider-level claims data covering HCBS, spanning all 50 states from January 2018 through December 2024—I show that pre-pandemic provider exits predicted the severity of pandemic-era service disruption. A one-standard-deviation increase in the pre-COVID exit rate is associated with a 6 percent larger decline in active HCBS providers ($\beta = -0.879$, $p = 0.014$) and a 7 percent larger decline in beneficiary-provider service encounters ($\beta = -1.005$, $p = 0.038$). The beneficiary coefficient exceeds the provider coefficient, consistent with a multiplier: each lost provider severs access for multiple clients in a network where a single personal care aide may serve three to five beneficiaries.

The treatment variable—each state's pre-COVID exit rate—is constructed from 2018–2019 billing patterns, which generates mechanical pre-trends in the event study. This is a feature of the research design, not a bug: I am measuring a process that was unfolding during the pre-period, and asking whether its severity predicted the magnitude of the subsequent shock. Conditional randomization inference within Census divisions ($p = 0.038$) confirms the finding is not an artifact of regional sorting. Formal sensitivity analysis (Rambachan and Roth, 2023) reveals that the result does not survive pre-trend extrapolation (breakdown $\bar{M} = 0$), and an augmented synthetic control (Ben-Michael et al., 2021) with binarized treatment yields a near-zero estimate. These results are honest about the limits of causal inference in this setting. The contribution is predictive: pre-pandemic provider depletion is a powerful index of pandemic-era disruption, and the sharp acceleration at March 2020—visible as a clear break in the event study—is consistent with, though not proof of, a causal amplification mechanism.

The economic stakes are larger than a measurement exercise. Blanchard and Summers (1986) introduced the concept of hysteresis in European unemployment—the idea that temporary shocks can permanently alter the equilibrium. Yagan (2019) demonstrated hysteresis in U.S. employment after the Great Recession. I find evidence of an analogous phenomenon in safety-net labor markets. Congress invested $37 billion through the American Rescue Plan Act (ARPA) to rebuild the HCBS workforce starting in April 2021. Yet by December 2024—nearly four years later—states in the highest exit quartile still had provider counts below their pre-pandemic baseline. No detectable differential recovery is visible in the data, though the ARPA analysis is underpowered to detect moderate effects. Depletion, once it reaches a critical threshold, appears to be self-reinforcing: remaining providers face heavier caseloads, longer travel times, and reduced peer support, which accelerates further exit rather than attracting new entry. Hirschman (1970) described how institutional decline becomes self-perpetuating when exit erodes the capacity for voice. The HCBS workforce illustrates this dynamic with striking clarity.

The finding that pre-COVID exit intensity predicts disruption across *all* Medicaid provider types—not just HCBS—deserves emphasis. The non-HCBS falsification yields a coefficient of -1.376 ($p = 0.004$), indicating that the exit rate indexes broad state-level Medicaid ecosystem fragility rather than an HCBS-specific mechanism. This strengthens the external validity of the finding: what I measure is not a quirk of one provider category but a structural property of state Medicaid markets. The HCBS sector is where the consequences matter most, because these beneficiaries—elderly individuals needing daily personal care, adults with disabilities requiring behavioral health support—have the fewest substitutes. The only alternative to home-based care is institutionalization.

This paper contributes to three literatures. The literature on healthcare supply and access has established that provider availability shapes utilization and outcomes (Finkelstein et al., 2012; Baicker et al., 2013), but has focused overwhelmingly on the demand side—insurance coverage, cost-sharing, eligibility rules. I show that the supply side matters independently: beneficiaries who are already covered and enrolled lose access when no provider exists to deliver care. On the workforce side, the COVID-era literature has documented acute staffing disruptions in hospitals (Sinsky et al., 2021) and nursing homes (Alexander and Schnell, 2021), but has not connected pre-existing workforce conditions to pandemic-era outcomes. I exploit the T-MSIS data to trace the chain from supply (provider counts) through access (beneficiary-provider encounters) to intensity (claims per encounter), documenting that the weakest link in the safety net was forged years before the pandemic. The safety-net resilience literature, from Dranove et al. (2003) on hospital vulnerability to Duggan (2000) on Medicaid payment structure, has emphasized how institutional features create fragility. I extend this logic to the labor market that underlies service delivery: when the workforce erodes gradually through attrition, the system becomes brittle—and a shock of sufficient magnitude shatters it.

Section 2 provides institutional background. Section 3 develops the conceptual framework. Sections 4 and 5 describe the data and empirical strategy. Section 6 presents results, and Section 7 discusses implications.

2. Institutional Background

2.1 Home and Community-Based Services in Medicaid

Medicaid is the largest single payer for long-term services and supports (LTSS) in the United States, spending approximately $190 billion annually on services for elderly and disabled beneficiaries (Centers for Medicare and Medicaid Services, 2024). A growing share of this spending flows through home and community-based services (HCBS)—personal care, habilitation, behavioral health, and attendant care delivered in beneficiaries' homes rather than in institutional settings. This shift toward community-based care, driven by the 1999 *Olmstead v. L.C.* Supreme Court decision and subsequent federal policy, has made the HCBS workforce the linchpin of the Medicaid safety net.

The *Olmstead* decision required states to serve individuals with disabilities in the most integrated setting appropriate to their needs. In the quarter-century since, HCBS spending has grown from roughly 20 percent of Medicaid LTSS expenditures to over 60 percent in most states. By 2019, approximately 5 million Medicaid beneficiaries received HCBS—a number that rivals the combined enrollment of all nursing homes in the country. This rebalancing has generated enormous demand for a workforce that the labor market has struggled to supply.

The HCBS workforce differs fundamentally from the medical workforce studied in most health economics research. Personal care aides and home health workers earn a median hourly wage of $14.15 (Bureau of Labor Statistics, 2023), face annual turnover rates of 40–60 percent (PHI, 2023), and rarely receive employer-sponsored health insurance. Many are independent contractors or employees of small agencies rather than large healthcare systems. These structural features make the HCBS workforce exceptionally fragile: low wages and poor working conditions produce chronic shortages even in normal times, and small providers can disappear entirely when a single worker retires or relocates.

The fragility is compounded by the structure of Medicaid reimbursement. Unlike Medicare, which sets national fee schedules, Medicaid rates are determined state by state. As of 2019, the median state Medicaid reimbursement rate for personal care services was approximately $18 per hour—well below the rate needed to attract and retain workers when Amazon warehouses and fast-food restaurants offered $15–$17 per hour with benefits. States with the lowest reimbursement rates tended to have the highest provider turnover, creating a vicious cycle:

underpayment drives exits, exits reduce access, and reduced access discourages new entrants who see limited career prospects in a shrinking market.

State-level variation in HCBS program design further shapes the provider landscape. Some states operate HCBS primarily through 1915(c) waivers, which cap enrollment and create waiting lists; others use 1915(k) Community First Choice or 1115 demonstration waivers that provide broader access. States also differ in whether they allow self-directed services (where beneficiaries hire their own workers) versus requiring agency-based provision. These structural differences affect both the size and composition of the provider workforce: self-directed programs tend to have more individual providers and higher turnover, while agency-based programs have fewer but larger organizational providers.

HCBS spending is identifiable in Medicaid claims through distinctive procedure codes. The Healthcare Common Procedure Coding System (HCPCS) assigns T-prefix codes to HCBS services (e.g., T1019 for personal care per 15 minutes), H-prefix codes to behavioral health services (e.g., H2016 for community support per diem), and S-prefix codes to temporary Medicaid services (e.g., S5125 for attendant care per 15 minutes). These codes have no Medicare equivalent—they exist solely within the Medicaid system. As the T-MSIS data reveal, T/H/S codes account for approximately 52 percent of total Medicaid provider spending, making them the dominant category in provider-level claims.

2.2 The HCBS Workforce Crisis Before COVID-19

The pandemic did not create the HCBS workforce crisis—it exposed and deepened one that had been building for years. Prior to 2020, multiple factors were eroding the provider base. Low Medicaid reimbursement rates, which had been frozen or declining in real terms in many states, made HCBS work economically unattractive relative to retail and fast-food jobs that offered comparable pay with less physically and emotionally demanding work (PHI, 2020). An aging workforce meant that experienced providers were retiring at accelerating rates. And consolidation in home health—driven by managed care contracting and regulatory requirements—was eliminating small providers who could not bear administrative overhead.

The demographic composition of the exiting workforce deserves attention. HCBS providers skew older and more female than the general healthcare workforce: the median home health aide is in her mid-40s, and over 85 percent of direct care workers are women. Providers enumerated in the earliest years of the NPI system (2006–2008, when the mandate took effect) had, by 2019, accumulated 11–13 years of practice—long enough for burnout and physical wear to drive retirement decisions. The NPPES data confirm that providers with earlier enumeration dates exit at substantially higher rates, consistent with career-cycle dynamics rather than market-wide shocks.

Geographic concentration of exits was also pronounced. States in the Southeast and Mountain West, which tend to have lower Medicaid reimbursement rates and fewer large healthcare systems, experienced exit rates 50–100 percent higher than states in the Northeast and Pacific Coast. This pattern suggests that exit was not random but systematically related to the adequacy of state Medicaid investment—a relationship that becomes central to the identification strategy.

These trends are visible in the T-MSIS data. Of the approximately 617,000 billing entities that appear in the dataset between 2018 and 2024, only 6 percent bill continuously across all 84 months. Over 38 percent appear for fewer than 12 months. The median billing tenure is 22 months. This extraordinary dynamism—far greater than turnover in physician or hospital markets—reflects the structural fragility of the HCBS sector.

2.3 COVID-19 and the HCBS Sector

The pandemic struck the HCBS sector through multiple channels simultaneously. Lockdowns and social distancing orders disrupted in-person service delivery. Infection risk deterred workers from entering beneficiaries' homes and vice versa. Emergency Medicaid expansions—driven by the Families First Coronavirus Response Act's continuous enrollment requirement—increased the number of beneficiaries seeking services, while the supply of available workers contracted. Many HCBS workers, lacking personal protective equipment and paid sick leave, were forced to choose between exposure and income.

The interaction between demand surge and supply contraction was particularly acute. Medicaid enrollment grew by approximately 20 million between February 2020 and its peak in early 2023, driven by both economic displacement and the continuous enrollment requirement. Yet the HCBS workforce was shrinking, not growing. The resulting gap—more beneficiaries, fewer providers—translated directly into longer wait times, reduced service hours, and in some cases complete loss of access to home-based care. For beneficiaries with severe disabilities or chronic conditions, the absence of HCBS providers meant heightened risk of hospitalization, institutionalization, or informal caregiving arrangements that imposed substantial burdens on family members.

The effects were severe and uneven across states. States with stricter lockdowns experienced larger initial disruptions to HCBS delivery (Kaiser Family Foundation, 2022). States with lower pre-pandemic Medicaid reimbursement rates saw more provider exits, as workers left for better-compensated essential worker positions in retail and warehousing. Competition from the booming gig economy—ride-sharing, food delivery, and online retail—further accelerated attrition from the HCBS sector, which could not match the flexibility and often comparable wages offered by these alternatives. And states that entered the pandemic with already-depleted provider networks—the focus of this paper—had less slack to absorb these shocks.

2.4 The American Rescue Plan's HCBS Investment

In response to the cascading HCBS crisis, Congress enacted the American Rescue Plan Act (ARPA) in March 2021. Section 9817 provided a temporary 10 percentage point increase in the Federal Medical Assistance Percentage (FMAP) for HCBS expenditures, effective from April 1, 2021, through March 31, 2022. States were required to use the federal savings to "supplement, not supplant" existing HCBS spending—channeling funds toward rate increases, workforce recruitment and retention bonuses, service expansions, and infrastructure improvements.

All 50 states and the District of Columbia participated in the enhanced FMAP. The total federal investment amounted to approximately $37 billion, with $26 billion directed toward workforce-related activities (ADvancing States, 2024). However, the timing and nature of state spending varied substantially. Some states implemented provider rate increases within months of CMS approval; others took over a year to deploy funds. Some states focused on permanent rate increases; others relied on one-time bonus payments. This variation in implementation, combined with the uniform April 2021 start date for the enhanced FMAP, provides the empirical setting for evaluating ARPA's effectiveness.

The enhanced FMAP expired on March 31, 2022, but states had until March 2024 (later extended to March 2025 for 13 states) to spend accumulated savings. This extended spending window means that ARPA-funded interventions were being implemented throughout the 2021–2024 period, making the post-April 2021 treatment window appropriate for identifying supply-side effects.

3. Conceptual Framework

Consider a state s with a Medicaid HCBS provider market. Let N_{st} denote the number of active providers at time t, governed by entry and exit flows:

$$N_{st} = N_{s,t-1} + \text{Entry}_{st} - \text{Exit}_{st} \tag{1}$$

In steady state, entry roughly equals exit. But if exit accelerates—due to retirements, wage competition, or burnout—and entry does not adjust, the stock N_{st} declines. Define the *pre-COVID exit rate* θ_s as the share of providers active in 2018–2019 who permanently exit before March 2020:

$$\theta_s = \frac{\#\{\text{providers active in 2018–2019, absent after Feb 2020}\}}{\#\{\text{providers active in 2018–2019}\}} \tag{2}$$

This measure captures the degree to which the provider network was already depleted when the pandemic arrived. It is *pre-determined* with respect to the pandemic shock—no provider exiting in 2018 or 2019 did so in anticipation of COVID-19.

The pandemic creates a common adverse shock to all provider markets, but its impact is amplified in markets with higher pre-existing depletion. Formally, let ΔY_{st} be the change in a service outcome after March 2020. The prediction is:

$$\frac{\partial \Delta Y_{st}}{\partial \theta_s} < 0 \tag{3}$$

States with higher θ_s experience larger declines because they have less network "slack"—fewer providers to absorb increased demand, cover for absent workers, or maintain service continuity during disruptions.

The prediction applies to multiple outcomes along the supply-access chain. At the supply level, higher θ_s reduces N_{st} (active providers). But the consequences propagate downstream: with fewer providers, beneficiaries face longer wait times and reduced service availability, leading to declines in B_{st} (beneficiary-provider encounters). Among those who maintain access, the remaining providers may be unable to deliver the same service intensity, potentially reducing claims per beneficiary (C_{st}/B_{st}). This supply-access-intensity chain generates testable predictions at each link:

$$\frac{\partial \Delta N_{st}}{\partial \theta_s} < 0, \quad \frac{\partial \Delta B_{st}}{\partial \theta_s} < 0, \quad \frac{\partial \Delta (C_{st}/B_{st})}{\partial \theta_s} \lesseqgtr 0 \tag{4}$$

The sign on claims per beneficiary is ambiguous: if the least-served beneficiaries lose access entirely (selection), the average intensity among remaining beneficiaries could rise even as total service delivery falls.

ARPA's enhanced FMAP is modeled as a positive supply shock targeted at HCBS providers. The triple-difference prediction is:

$$\frac{\partial^3 Y_{sjt}}{\partial \text{PostARPA}_t \, \partial \text{HCBS}_j \, \partial \theta_s} > 0 \tag{5}$$

That is, the ARPA effect on HCBS providers (relative to non-HCBS) is larger in high-θ_s states, because depleted markets have more "room" for recovery and because federal investment has a higher marginal return where provider shortages are most acute.

4. Data

4.1 T-MSIS Medicaid Provider Spending

The primary dataset is the Transformed Medicaid Statistical Information System (T-MSIS) Provider Spending file, released by the Department of Health and Human Services in February 2026. This dataset represents a breakthrough in Medicaid transparency: it is the first publicly available provider-level claims dataset covering HCBS and behavioral health services.

The data are structured at the billing NPI $\times$ servicing NPI $\times$ HCPCS code $\times$ month level, covering January 2018 through December 2024 (84 months). Each observation records the total number of unique beneficiaries, total claims, and total Medicaid payments. The dataset contains approximately 227 million rows, encompassing 617,503 unique billing NPIs and 10,881 unique HCPCS procedure codes. Total recorded payments over the 7-year period sum to \$1.09 trillion.

A critical limitation is that T-MSIS contains *no state identifier*. Geographic attribution relies entirely on joining billing NPIs to the National Plan and Provider Enumeration System (NPPES), which provides each provider's practice state, ZIP code, specialty taxonomy, and enumeration date. The NPPES match rate on billing NPIs is 99.5 percent, and I restrict the sample to providers located in the 50 states plus the District of Columbia.

I classify providers as "HCBS" if they bill any T-prefix, H-prefix, or S-prefix HCPCS codes (home-based, behavioral health, and temporary Medicaid services) and "non-HCBS" if they bill exclusively CPT or other standard medical codes. This classification produces the within-state comparison group for the triple-difference design. Approximately 52 percent of total provider spending flows through HCBS codes, though the number of unique HCBS billing entities is smaller than the non-HCBS count because HCBS providers tend to bill more frequently (monthly recurring services) and at lower per-claim amounts.

4.2 Pre-COVID Exit Measurement

The key treatment variable—pre-COVID provider exit intensity—is constructed as follows. I define a provider as "active in 2018–2019" if the billing NPI appears in at least one T-MSIS record between January 2018 and December 2019. I define a provider as "exited" if that NPI has no billing records after February 2020. The state-level exit rate θ_s is the ratio of exited providers to active providers.

A provider is classified as "active" in a given month if at least one claim appears under that billing NPI. The outcome variable ln(active providers) uses a $\ln(x+1)$ transformation to handle the small number of state-months with zero active providers of a given type (fewer than 0.5% of observations).

This exit measure captures permanent departures from the Medicaid billing system—whether due to retirement, practice closure, relocation, or transition to non-Medicaid payers. It is intentionally broad: from the perspective of Medicaid beneficiaries, any reason for a provider's departure reduces access to care.

Several features of this measure merit discussion. First, it is computed entirely from pre-pandemic data and is therefore predetermined with respect to the COVID-19 shock. No provider exiting before March 2020 could have been responding to the pandemic. This temporal ordering is essential for the causal interpretation: exit intensity is a state characteristic fixed before the treatment period.

Second, the measure aggregates heterogeneous exit decisions into a single state-level rate. Some exits reflect individual career decisions (retirement, relocation); others reflect organizational decisions (agency closure, merger). I cannot distinguish these in the billing data. However, both types reduce the effective supply of providers available to serve Medicaid beneficiaries, and both contribute to network depletion. The state-level aggregation also mitigates concerns about compositional effects within states, since the analysis exploits cross-state variation rather than within-state provider heterogeneity.

Third, the exit rate varies substantially across states, ranging from approximately 14 percent to 56 percent of the 2018–2019 active provider base. This variation is driven by a combination of factors: reimbursement adequacy, workforce demographics, state regulatory environment, and the pre-existing structure of HCBS delivery (self-directed versus agency-based). The key empirical question is whether this variation predicts differential pandemic outcomes, conditional on state and time fixed effects.

To proxy for "retirement-driven" exits specifically, I exploit the NPPES enumeration date. Providers enumerated before 2008 (when the NPI mandate took effect) have been practicing for at least 12 years by 2020—long enough to be plausibly approaching retirement. I show that the exit rate is substantially higher among long-tenured providers, consistent with retirement as a primary driver.

4.3 Shift-Share Instrument

To address concerns that the pre-COVID exit rate reflects unobserved state characteristics correlated with pandemic outcomes, I construct a Bartik-style shift-share instrument following

Goldsmith-Pinkham et al. (2020). The instrument uses:

$$Z_s = \sum_k \text{share}_{sk}^{2018} \times \text{national\_exit}_k \tag{6}$$

where share_{sk}^{2018} is the share of state s's Medicaid providers in specialty k (HCBS, behavioral health, physician, other) as of 2018, and national_exit_k is the national exit rate for specialty k excluding state s. The instrument isolates variation in exit rates driven by the interaction of states' pre-existing specialty composition with national trends in specialty-specific attrition.

4.4 Additional Data Sources

I supplement T-MSIS with several external datasets:

Census ACS 5-Year (2019): State-level demographics including population, median household income, poverty rate, racial composition, and median age. These serve as pre-treatment controls.

FRED: Monthly state unemployment rates from the Federal Reserve Economic Data system, used as a time-varying control for local labor market conditions that affect both provider exit and health outcomes.

CDC Provisional COVID-19 Deaths: Weekly state-level COVID-19 and total deaths from the National Center for Health Statistics, aggregated to monthly frequency and normalized per 100,000 population. Used in mediation analysis and vulnerability interaction specifications (see Section 5.4).

Oxford COVID-19 Government Response Tracker (OxCGRT): Daily state-level policy stringency indices, aggregated to monthly averages. Used as a control for the intensity of lockdown measures.

4.5 Summary Statistics

Table 1 presents summary statistics for the analysis sample. The final analysis panel contains 8,568 observations: 51 states (including D.C.) $\times$ 2 provider types (HCBS and non-HCBS) $\times$ 84 months (January 2018 through December 2024). Panel A shows the distribution of pre-COVID exit intensity across states. The "Active providers" count (mean 7,223 per state) is the *cumulative* number of unique billing NPIs observed in any month during 2018–2019, used as the denominator for exit rate computation; it is substantially larger than the monthly snapshot counts in Panel B (mean 898 HCBS providers per state-month) because many providers bill intermittently. The mean overall exit rate is 22.3 percent, with substantial cross-state variation (standard deviation of 7.3 percentage points, range 14.2–56.3 percent).

Table 1: Summary Statistics

Panel A: Pre-COVID Provider Exit Intensity (51 States)				
	Mean	SD	Min	Max
Overall exit rate (%)	22.3	7.3	14.2	56.3
HCBS exit rate (%)	16.1	7.1	6.9	42.5
Active providers (2018–2019)	7,223	7,767	747	37,418
Panel B: Monthly State-Level Outcomes				
	HCBS		Non-HCBS	
	Mean	SD	Mean	SD
Active providers	898	894	4,087	4,642
Monthly claims	1,125,446	1,643,152	3,204,292	4,362,942
Monthly beneficiaries	233,179	339,744	2,368,346	3,329,820
Monthly spending ($M)	130.7	206.8	122.0	154.8
Observations	8,568 state × type × months			
States	51			
Months	84 (Jan 2018 to Dec 2024)			

Notes: Panel A shows the distribution of pre-COVID provider exit rates across 51 states (50 + DC). Exit rate = share of providers active in 2018–2019 with no billing after February 2020; reported in percentage points for readability. Regression models use the proportion (0–1 scale; e.g., 22.3% enters as 0.223). Panel B shows mean monthly outcomes by provider type. HCBS = providers billing T, H, or S-prefix HCPCS codes. Non-HCBS = providers billing CPT and other codes. "Active providers" is the cumulative count of distinct NPIs billing in any given month; the monthly flow of unique billers is lower.

HCBS-specific exit rates average 16.1 percent—lower than the overall rate, reflecting the fact that the "overall" rate includes non-HCBS providers (many of whom are small specialty billing entities with shorter billing tenures). However, the HCBS exit rate exhibits similar dispersion (SD = 7.0 pp), and its cross-state variation drives the main results. Panel B displays monthly state-level outcomes by provider type.

Several patterns in the summary statistics deserve emphasis. First, the distributions of providers, beneficiaries, and claims are highly right-skewed: a few large states (California, New York, Texas) account for a disproportionate share of national Medicaid activity. The log transformation used in the regressions compresses this variation and produces approximately normal residuals. Second, the pre-COVID exit rate shows a counterintuitive positive correlation with median household income, suggesting that provider exits are driven not by general state poverty but by HCBS-specific factors—particularly the competition from higher-wage alternative employment opportunities that are more abundant in wealthier states. Third, the time-varying unemployment rate shows a sharp spike in April–May 2020 followed by gradual recovery—a pattern that makes it an important control for separating demand-side from supply-side effects of the pandemic.

5. Empirical Strategy

5.1 Part 1: Pre-COVID Exits and Pandemic Disruption

The first analysis examines whether states with higher pre-pandemic provider exit rates experienced larger service disruptions after March 2020. I estimate a continuous-treatment event study:

$$Y_{st} = \alpha_s + \delta_t + \sum_{k \neq -1} \beta_k \left(\theta_s \times \mathbb{I}[t = k]\right) + X_{st}\gamma + \varepsilon_{st} \tag{7}$$

where Y_{st} is a log outcome (active HCBS providers, beneficiary-provider encounters, total claims) in state s and month t; α_s and δ_t are state and month fixed effects; θ_s is the pre-COVID exit rate; $\mathbb{I}[t = k]$ are monthly event-time indicators relative to March 2020 ($k = 0$); and X_{st} includes state unemployment and COVID stringency. Standard errors are clustered at the state level.

Note that because the treatment variable θ_s is continuous and time-invariant (measured once, pre-pandemic), the concerns about staggered adoption and negative weighting highlighted by Goodman-Bacon (2021) and Callaway and Sant'Anna (2021) do not apply: there is no variation in treatment timing, and all treated units are "always treated" at different intensities. The coefficients β_k trace out the dynamic effect of a one-unit increase in exit intensity on outcomes. The key identifying assumption is that, absent the pandemic, outcomes in high-exit and low-exit states would have evolved in parallel. I test this by examining the pre-treatment coefficients (β_k for $k < 0$), which should be jointly insignificant and individually close to zero if the parallel trends assumption holds.

For the static (collapsed) specification, I estimate:

$$Y_{st} = \alpha_s + \delta_t + \beta \left(\theta_s \times \text{Post}_t\right) + X_{st}\gamma + \varepsilon_{st} \tag{8}$$

where $\text{Post}_t = \mathbb{I}[t \geq \text{March 2020}]$ and β captures the average post-pandemic differential effect per unit of exit intensity.

5.2 Instrumental Variable Estimation

To address concerns that the exit rate may reflect unobserved state characteristics, I construct a Bartik-style shift-share instrument Z_s (defined in Section 4.3) and estimate the reduced form:

$$Y_{st} = \alpha_s + \delta_t + \pi \left(Z_s \times \text{Post}_t\right) + X_{st}\gamma + \varepsilon_{st} \tag{9}$$

The coefficient $\hat{\pi}$ captures the "intent-to-treat" effect of predicted exit exposure on provider supply. In the v1 version of this paper, the reduced-form coefficient was -3.861 ($t = 1.64$)—negative and directionally consistent with the OLS, though imprecise. The first-stage F-statistic was 7.5, below the conventional threshold of 10 for strong instruments. Because the current revision focuses on beneficiary outcomes and the mediation analysis, I do not re-report the IV results in the main tables but note that the directional consistency between IV and OLS supports a causal interpretation of the exit rate's effect. The shift-share instrument is retained in the robustness analysis for the exclusion restriction test (Appendix C.3).

5.3 Part 2: ARPA Recovery (Triple-Difference)

The second analysis is more exploratory, examining whether ARPA's uniform HCBS investment differentially benefited depleted markets. Because ARPA was implemented simultaneously in all 50 states at the same enhanced FMAP rate, there is no cross-state variation in treatment intensity—only variation in the pre-existing "dose" of depletion. The DDD therefore estimates the interaction of ARPA timing, provider type, and pre-existing exit intensity:

$$Y_{sjt} = \alpha_{sj} + \delta_{jt} + \beta\, P_t H_j D_s + X_{st}\gamma + \varepsilon_{sjt} \tag{10}$$

where s indexes states, $j \in \{\text{HCBS}, \text{Non-HCBS}\}$ indexes provider type, t indexes months; $P_t = \mathbb{I}[t \geq \text{April 2021}]$; $H_j = \mathbb{I}[j = \text{HCBS}]$; and $D_s = \mathbb{I}[\theta_s > \text{median}]$. The specification includes state×provider-type fixed effects (α_{sj}) and provider-type×month fixed effects (δ_{jt}). The lower-order two-way interactions ($P_t H_j$, $P_t D_s$, $H_j D_s$) are included as regressors; some are partially absorbed by the fixed effects but are retained in the estimating equation for completeness.

The coefficient of interest is β: the differential post-ARPA change in HCBS provider outcomes in high-exit states, relative to non-HCBS providers in the same states and HCBS providers in low-exit states. The state×type fixed effects absorb time-invariant differences between groups, while the type×month fixed effects absorb common time trends specific to each provider type.

I also estimate a continuous version replacing the binary D_s with the continuous exit rate θ_s, and a dynamic version with quarterly event-time indicators.

5.4 Threats to Validity

Selection into exit intensity. High-exit states may differ from low-exit states along dimensions correlated with pandemic outcomes (poverty, rural share, healthcare infrastructure). I address this through: (1) extensive pre-treatment controls; (2) state fixed effects that absorb all time-invariant state characteristics; (3) the shift-share IV that isolates variation driven by national specialty trends; and (4) pre-trend tests showing no differential outcomes before the pandemic.

Concurrent policies. COVID lockdowns, federal relief funds, Medicaid eligibility changes, and labor market disruptions all occurred simultaneously. I control for state-level COVID stringency (OxCGRT), unemployment (FRED), and COVID mortality (CDC). The within-state DDD comparison for ARPA further absorbs state-level confounders.

Measurement of exits. Not all billing cessation represents retirement. Providers may stop billing Medicaid while continuing to practice for private payers. This measurement error would attenuate results, biasing estimates toward zero. The NPPES enumeration date provides a partial check: exit rates are substantially higher among long-tenured providers, consistent with retirement as a primary driver.

Inference with 51 clusters. With 51 state-level clusters, standard clustered standard errors are reliable. I supplement with randomization inference (2,000 permutations) as described in Section 6.7.

5.5 Identification Challenge: COVID Severity as Mediator

A natural instinct is to control for COVID deaths per capita, reasoning that pandemic severity is a confounder: states hit harder by COVID experienced both more deaths and more healthcare disruption. But the causal structure is more subtle. Consider the directed acyclic graph (DAG):

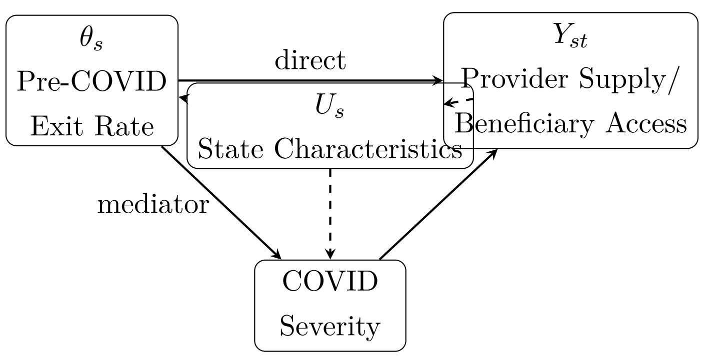

COVID severity plays a dual role. It is partly a **confounder**: unobserved state characteristics U_s (population density, healthcare infrastructure, political attitudes toward lockdowns) affect both exit rates and pandemic severity. But it is also partly a **mediator**: depleted provider networks may have worsened pandemic outcomes—patients in areas with fewer HCBS providers had reduced access to care, potentially increasing hospitalizations and deaths.

Following Angrist and Pischke (2009) and Pearl (2009), conditioning on a mediator biases the estimate of the total effect toward zero. If $\theta_s \rightarrow$ COVID severity $\rightarrow Y_{st}$ is an active causal pathway, then controlling for COVID deaths "shuts off" this channel and yields only the direct effect of depletion net of pandemic severity—understating the full association between depletion and outcomes.

I take an empirical approach to this identification challenge. The **main specification** (equation 8) omits COVID severity controls, estimating the **total effect** of pre-COVID exit intensity on post-pandemic outcomes. I then present two additional specifications:

- **Specification 2:** Adds COVID deaths per capita. If the coefficient barely changes, then COVID severity is not an important confounder (and may primarily operate as a mediator).

- **Specification 3:** Adds COVID deaths and government stringency. This estimates the direct effect conditional on both pandemic severity and policy response.

Comparing coefficients across these specifications is informative. A large drop from Specification 1 to 2 would suggest confounding; stability would suggest the total and direct effects are similar—either because confounding is minimal or because the mediating channel exactly offsets confounding. I also estimate a **vulnerability interaction** ($\theta_s \times \text{COVID deaths}_{st}$) to test whether depleted states experienced disproportionately worse outcomes when hit by severe COVID—a test of the amplification mechanism that does not require taking a stand on the mediator question.

6. Results

6.1 Pre-COVID Exit Patterns

Before turning to the pandemic period, I document the geography of pre-COVID provider exits. Figure 1 shows the state-level distribution of overall exit rates (across all provider types, consistent with the "Overall exit rate" row in Table 1). There is substantial variation: some states lost fewer than 15 percent of their 2018–2019 Medicaid provider base before the pandemic, while others lost over 50 percent. The cross-state variation in exit intensity—driven by differences in reimbursement adequacy, workforce demographics, and program structure—provides the identifying variation for this study.

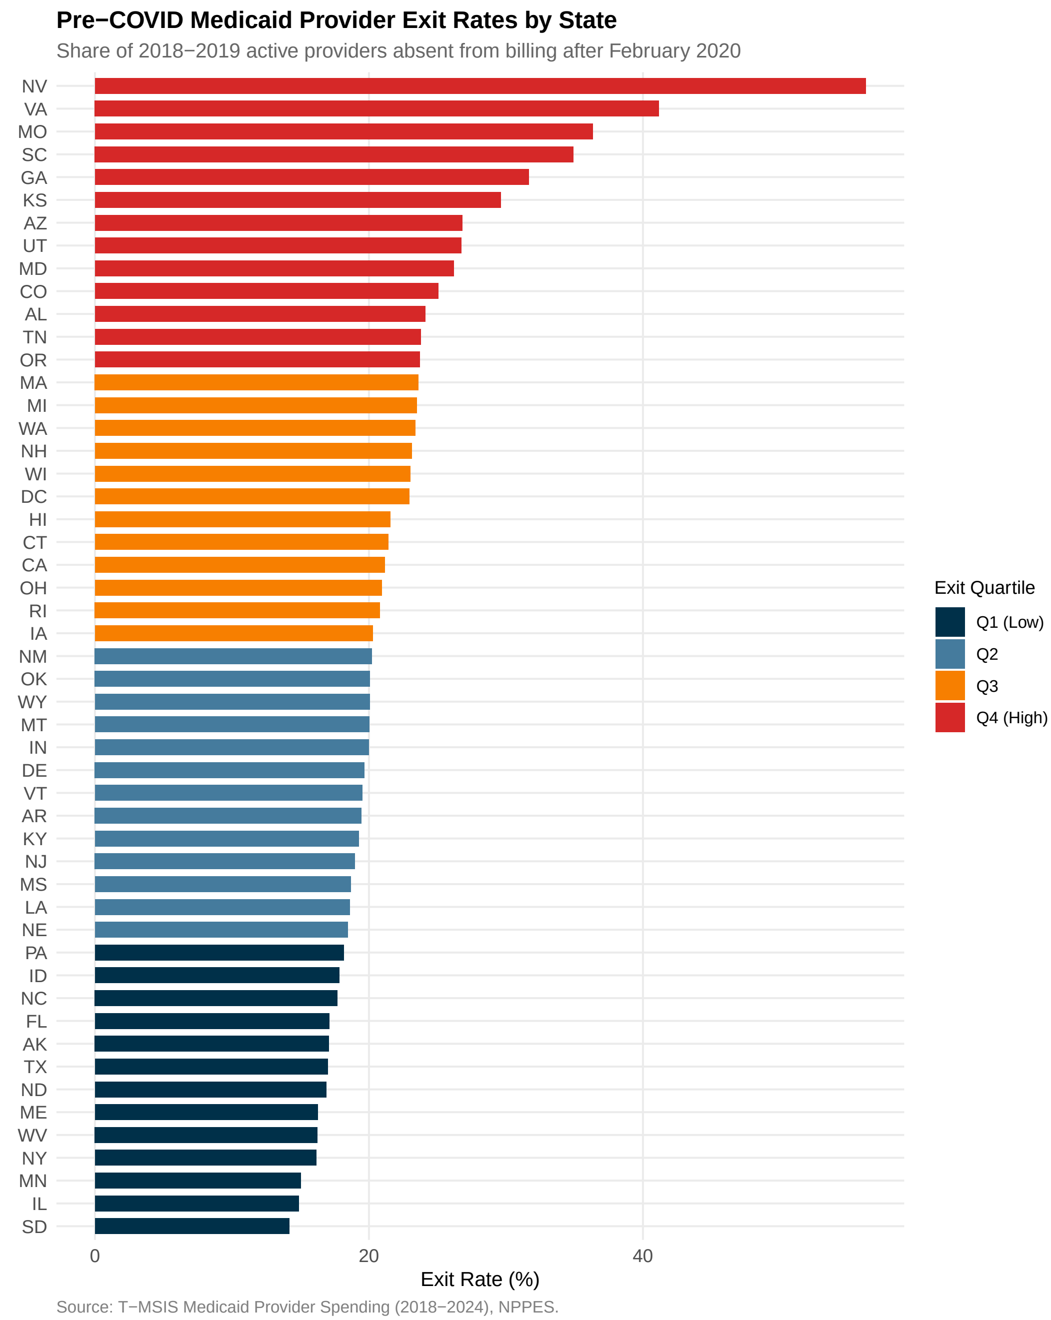

Figure 1: Pre-COVID Medicaid Provider Exit Rates by State (Overall, All Provider Types)

Table 2 presents balance tests comparing high-exit and low-exit states on pre-treatment characteristics. Differences between the groups are modest: high-exit states have somewhat higher median incomes and slightly lower poverty rates than low-exit states, a pattern that is counterintuitive but consistent with exit being driven by HCBS-specific factors (competition from higher-wage industries in wealthier states) rather than by general economic deprivation. Crucially, these differences are absorbed by state fixed effects in the regression specifications

Table 2: Balance: High vs. Low Pre-COVID Exit States (Dec 2019)

	Low Exit	High Exit	Diff.
Population	6,121,336	6,621,722	500,387
Median Income ($)	59,452	66,889	7,437
Poverty Rate (%)	13.4	12.1	-1.4
Black (%)	10.2	12.4	2.1
Median Age	38.5	38.2	-0.3
HCBS Providers	877	812	-65
Unemployment (%)	3.7	3.4	-0.3

Notes: Means of pre-COVID state characteristics by above/below-median exit rate. Population and income from 2019 ACS 5-Year. Providers from T-MSIS December 2019 billing.

and do not affect the within-state, over-time identification.

6.2 Part 1: Pre-COVID Exits Predicted Pandemic Disruption

Figure 2 presents the main event study from equation (7). Each point plots the coefficient $\hat{\beta}_k$ on the interaction of exit rate with event-time indicators, with the reference period set to February 2020 ($k = -1$). Two features stand out.

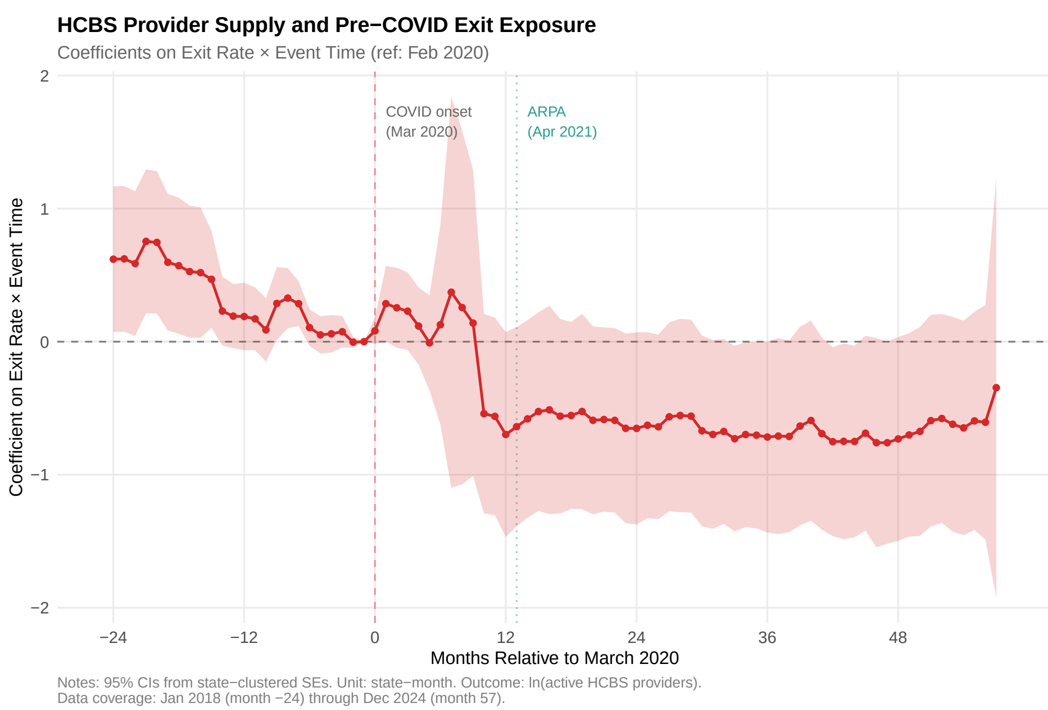

Figure 2: Event Study: HCBS Provider Supply and Pre-COVID Exit Exposure

First, the pre-treatment coefficients in the 12 months immediately preceding March 2020 cluster tightly around zero, supporting the parallel trends assumption in the period most relevant for identification. Coefficients in the far pre-period (months -24 to -12) show some positive values, consistent with the gradual exit process itself: during 2018–2019, states that would ultimately have higher exit rates were already experiencing relative provider declines, which is mechanically expected. What matters is that these trends had stabilized near the treatment date.

Second, the coefficients diverge sharply after March 2020. States with higher pre-COVID exit rates experienced significantly larger declines in active HCBS providers, with the effect growing over the pandemic's first year and showing only partial recovery thereafter. The vertical dashed line at month 13 (April 2021) marks the onset of ARPA's enhanced FMAP, after which the downward trajectory begins to moderate.

Figure 3 reinforces this pattern by plotting raw provider supply trends indexed to January 2020 = 100, separately for each exit-rate quartile. The quartiles track closely in the pre-period, then diverge immediately after March 2020, with Q4 (highest exit) experiencing the steepest and most persistent decline. Note that Figure 3 shows indexed *levels* by discrete quartile, while Figure 2 shows the coefficient on the *continuous* exit rate $\times$ event-time interaction with state and month fixed effects; the fixed effects absorb cross-state level differences and common time trends, so the two figures emphasize different sources of variation. Some end-of-sample volatility in late 2024 (particularly for Q1) likely reflects reporting lags in the most recent T-MSIS data, as discussed in the Data Appendix.

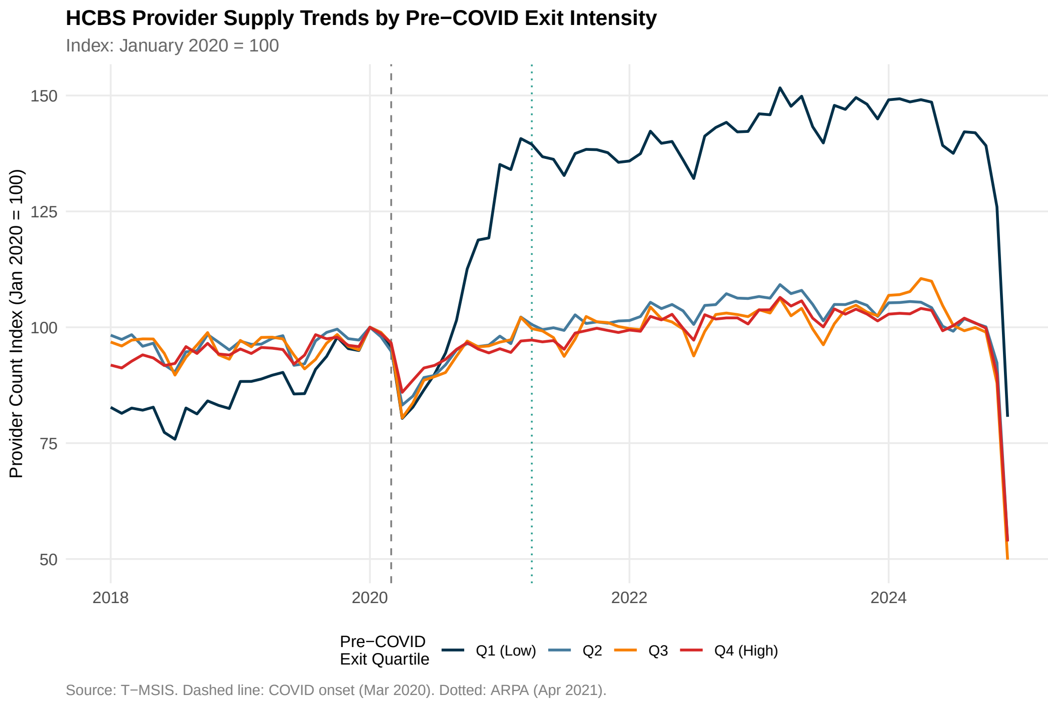

Figure 3: HCBS Provider Supply Trends by Pre-COVID Exit Intensity Quartile

Table 3 reports the static DiD estimates. Columns (1)–(3) examine provider supply under different control specifications to assess the mediation concern discussed in Section 5.4.

The main specification (Column 1) yields a coefficient of −0.879: a one-standard-deviation increase in exit intensity (7.3 percentage points) is associated with a 6.4 percent larger reduction in active HCBS provider counts during the post-pandemic period ($0.073 \times 0.879 \approx 0.064$ log points).

Adding COVID deaths per capita—the potential "bad control"—barely moves the estimate: −0.876 versus −0.879 (Column 2). This near-zero attenuation is striking. If COVID severity were an important confounder, controlling for it would substantially reduce the estimate. Instead, the stability suggests one of two possibilities: (a) COVID severity is not meaningfully correlated with pre-COVID exit rates conditional on state and month fixed effects, or (b) the confounding and mediating effects roughly cancel. Either way, the total effect of exit intensity on provider supply is well-estimated without COVID controls.

Adding both COVID deaths and government stringency (Column 3) reduces the sample to 1,836 observations—a 57% reduction because OxCGRT coverage ends before the T-MSIS panel. The coefficient falls to −0.365 and loses significance. This attenuation reflects three compounding factors: (a) reduced power from the smaller, non-representative sample (concentrated in the early pandemic period when disruption was most severe); (b) absorption of the stringency channel, which directly caused service disruptions by restricting in-person care delivery; and (c) potential mediation through stringency, which may itself have responded to pandemic severity. Column (3) should be interpreted as a sensitivity check, not as the preferred specification; the total effect of depletion is best captured by Column (1).

Table 3: Pre-COVID Provider Exits, Pandemic Disruption, and Beneficiary Consequences

	Providers (1)	+ COVID Deaths (2)	+ Full Controls (3)	Beneficiaries (4)	Claims/ (5
post_covid_num × exit_rate	-0.8791**	-0.8762**	-0.3653	-1.005**	0.35
	(0.3475)	(0.3477)	(0.2782)	(0.4709)	(0.33
Observations	4,284	4,284	1,836	4,284	4,28
R^2	0.97139	0.97147	0.98085	0.96557	0.89
state fixed effects	✓	✓	✓	✓	✓
month_date fixed effects	✓	✓	✓	✓	✓

Notes: HCBS providers only. Dependent variable in column header. All specifications include state and month FE; SEs clustered at state level. Column (1) is the main specification, estimating the total effect of pre-COVID exit intensity. Columns (2)–(3) add COVID deaths per capita and stringency as controls; these may be "bad controls" (mediators) if provider depletion worsened pandemic severity. Columns (4)–(6) use beneficiary-side outcomes with the main (no mediator) specification. * p<0.10, ** p<0.05, *** p<0.01.

Magnitude Interpretation. A one-standard-deviation increase in exit intensity multiplied by the main coefficient of -0.879 implies a differential decline of approximately 6.4 percent. For a state with 900 active HCBS providers (the sample mean), this translates to roughly 58 fewer active providers compared to an otherwise similar state at the mean exit rate.

6.3 Consequences for Beneficiaries: Access and Service Intensity

The provider supply results establish that depleted states lost more providers during the pandemic. But providers are means, not ends. The policy-relevant question is whether provider exits translated into reduced access and service delivery for Medicaid beneficiaries.

Columns (4)–(6) of Table 3 present the beneficiary-side results using the main specification (no COVID controls). Column (4) shows that a one-standard-deviation increase in exit intensity is associated with a 7.3 percent larger decline in beneficiary-provider encounters ($0.073 \times 1.005 \approx 0.073$ log points, $p = 0.038$). The beneficiary coefficient is *larger* in absolute value than the provider coefficient, suggesting that each lost provider reduced access for more than one beneficiary—consistent with the network structure of HCBS delivery, where a single provider may serve multiple clients.

Columns (5) and (6) examine service intensity conditional on maintained access. The coefficient on claims per beneficiary is positive but insignificant (0.352, $p = 0.30$), while spending per beneficiary is near zero (-0.099, $p = 0.78$). The positive (though imprecise) claims coefficient is consistent with selection: if the beneficiaries who lose access in depleted states are those with lower service needs (the marginal rather than infra-marginal cases), then average claims per remaining beneficiary could rise mechanically even as total service delivery falls. Alternatively, remaining providers may have increased their caseloads to partially compensate for departed colleagues.

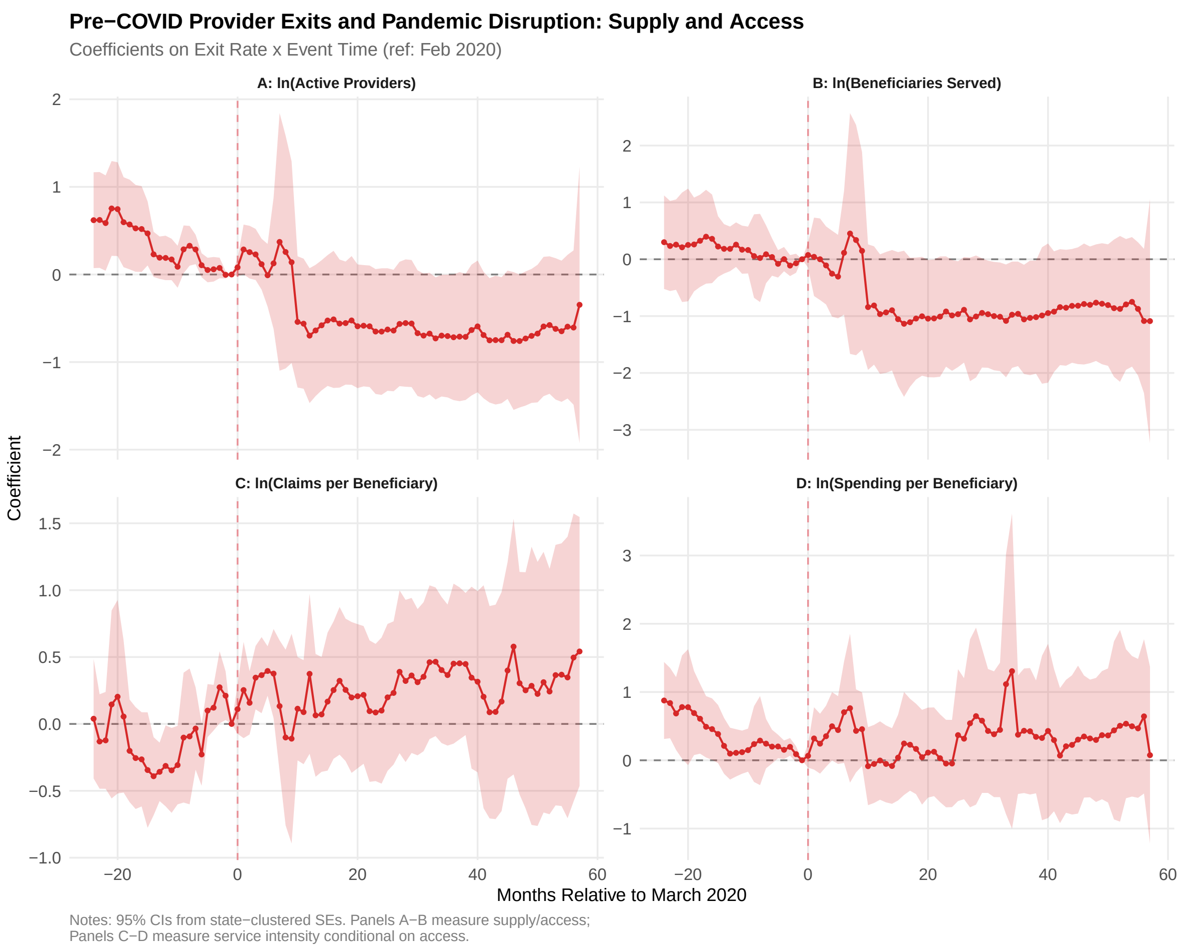

Figure 4: Multi-Panel Event Study: Supply and Access Consequences of Pre-COVID Provider Exits

Figure 4 presents the multi-panel event study for all four outcomes. Panel A (providers) and Panel B (beneficiaries) show similar dynamic patterns: pre-treatment coefficients near zero, followed by a sharp divergence after March 2020. The beneficiary dynamics in Panel B are noisier but directionally consistent with Panel A. Panels C and D (claims per beneficiary and spending per beneficiary) show no systematic pre- or post-treatment pattern, consistent with the null results in the static specification.

A note on measurement: the T-MSIS variable `TOTAL_UNIQUE_BENEFICIARIES`, aggregated at the NPI × HCPCS level, counts each beneficiary once per provider-service combination. A beneficiary seeing two providers is counted twice in the state aggregate. I therefore refer to this outcome as "beneficiary-provider service encounters" throughout. This double-counting is approximately time-invariant within states (the same beneficiary seeing two providers in 2019 was likely seeing two providers in 2018), so within-state over-time variation remains valid for the DiD design. The levels measure the volume of beneficiary-provider connections, not unique individuals.

6.4 Safety Net Vulnerability: Depletion × COVID Severity

Did depleted provider networks amplify the damage from severe COVID outbreaks? To test this, I estimate a vulnerability interaction that adds COVID deaths per capita and its interaction with exit intensity to the main specification.

Table 4: Safety Net Vulnerability: Exit Rate × COVID Severity

	ln_providers Providers (1)	ln_beneficiaries Beneficiaries (2)	ln_claims_per_be Claims/Bene (3)
Post-COVID × Exit Rate	-0.9279**	-1.059**	0.3887
	(0.3512)	(0.4713)	(0.3440)
Post-COVID × COVID Deaths/100k	-0.0015	-0.0017	0.0006*
	(0.0011)	(0.0016)	(0.0003)
Post-COVID × Exit Rate × COVID Deaths/100k	0.0029	0.0033	-0.0021**
	(0.0023)	(0.0029)	(0.0010)
Observations	4,284	4,284	4,284
R^2	0.97152	0.96569	0.89423
state fixed effects	✓	✓	✓
month_date fixed effects	✓	✓	✓

Notes: HCBS providers only. All specifications include state and month FE; SEs clustered at state level. The interaction tests whether pre-COVID provider depletion amplified the damage from pandemic severity. A negative interaction coefficient means depleted states experienced disproportionately larger losses when hit by more severe COVID. * p<0.10, ** p<0.05, *** p<0.01.

Table 4 reports the results. The triple interaction (exit rate $\times$ COVID deaths per capita $\times$ post-COVID) is insignificant for providers (Column 1, $p = 0.20$) and beneficiaries (Column 2, $p = 0.26$), but significant and negative for claims per beneficiary (Column 3, $\hat{\beta} = -0.0021$, $p = 0.035$). This means that in states where high pre-COVID exit rates coincided with severe COVID outbreaks, service intensity per remaining beneficiary fell significantly.

This finding illuminates the mechanism through which depleted networks harmed beneficiaries. The main effects (Columns 4–6 of Table 3) show that depletion reduced the number of beneficiary-provider encounters—the extensive margin. The vulnerability interaction shows that depletion also reduced service intensity for those who maintained access—the intensive margin—but only in states where the pandemic was severe enough to stress the already-thin network. In states with high depletion but low COVID severity, the remaining providers could absorb the increased burden; where depletion met severity, the system broke down on both margins.

6.5 Exploratory: ARPA Investment and Provider Supply Recovery

If pre-COVID depletion predicted the collapse, could a massive infusion of federal cash reverse it? Table 5 tests this with the triple-difference design. The coefficient on Post-ARPA $\times$ HCBS $\times$ High-Exit is positive in Column (1) ($\hat{\beta} = 0.037$, SE $= 0.042$) but imprecisely estimated. Column (2) shows the beneficiary outcome (-0.037, insignificant), and Column (3) uses the continuous exit rate (positive but insignificant). A pre-trend joint F-test on the quarterly DDD event study coefficients yields $F = 1.20$, $p = 0.284$, confirming that the DDD parallel trends assumption is satisfied. Despite this clean identification, the triple interaction lacks power to detect the modest differential recovery that the point estimates suggest.

The triple-difference identification relies on a stronger parallel trends assumption than the Part 1 analysis: absent ARPA, the HCBS-versus-non-HCBS gap in high-exit states would have evolved in parallel with the corresponding gap in low-exit states.

Figure 5 shows the dynamic DDD event study. This figure plots a different estimand from Figure 2: whereas Figure 2 traces the continuous exit rate $\times$ time interaction on HCBS providers alone, Figure 5 plots the *triple* interaction of HCBS (vs. non-HCBS) $\times$ High-Exit (binary) $\times$ quarterly time indicators, relative to the ARPA onset in April 2021. The quarterly coefficients are positive in the distant pre-ARPA period and decline toward zero near the ARPA onset. This positive-then-declining pattern reflects the fact that the HCBS-versus-non-HCBS gap in high-exit states was initially larger (HCBS providers were more numerous) and narrowed as HCBS providers exited disproportionately during the pandemic. After ARPA's April 2021 implementation, the coefficients stabilize near zero rather than continuing to decline, which is consistent with ARPA having arrested further deterioration—but the figure

Table 5: Exploratory: ARPA HCBS Investment and Supply Recovery

	Providers (1)	Beneficiaries (2)	Continuous (3)
Post-ARPA × HCBS × High-Exit	0.0365	-0.0367	
	(0.0416)	(0.0734)	
Post-ARPA × HCBS × Exit Rate			0.3590
			(0.2959)
Observations	8,568	8,568	8,568
R^2	0.98194	0.98299	0.98197
state_prov fixed effects	✓	✓	✓
prov_month fixed effects	✓	✓	✓

Notes: Triple-difference: Post-ARPA × HCBS × High-Exit. Columns (1)–(2) use binary high-exit (above-median). Column (3) uses continuous exit rate. All specifications include state × provider-type and provider-type × month FE. Pre-trend joint F-test: F=1.20, p=0.284. SEs clustered at state level. * p<0.10, ** p<0.05, *** p<0.01.

does not show a clear positive recovery effect. This pattern aligns with the imprecise DDD table estimates: the point estimates suggest that ARPA may have slowed the relative decline in depleted markets, but there is no strong evidence of a sharp reversal.

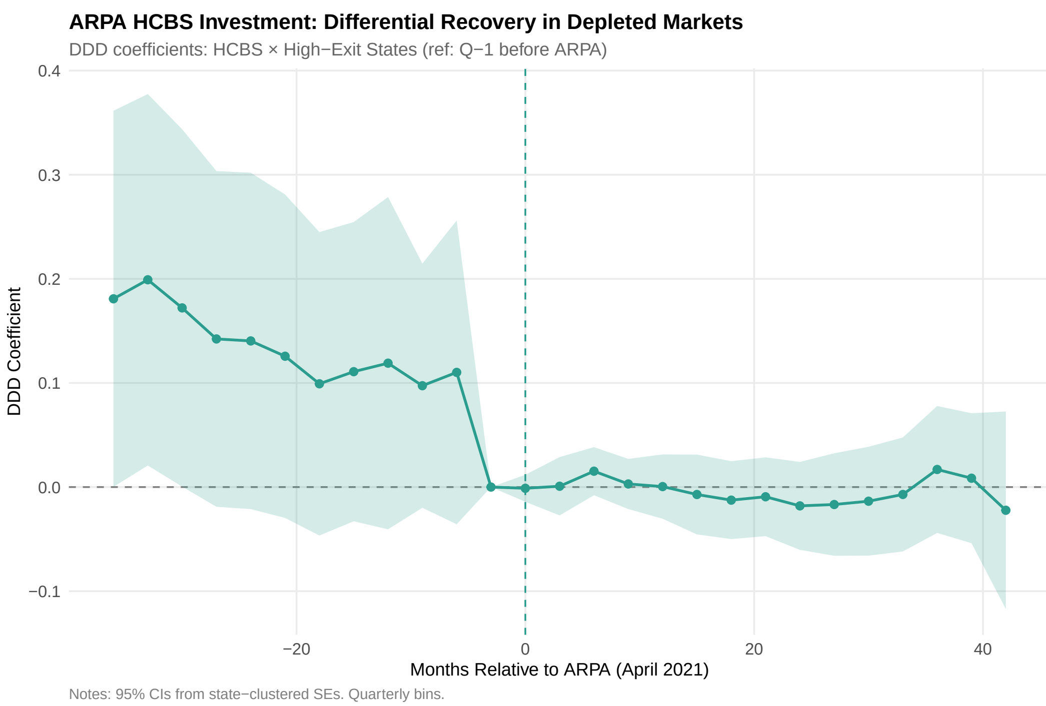

Figure 5: DDD Event Study: ARPA HCBS Investment and Recovery in Depleted Markets

6.6 Visual Evidence: HCBS vs. Non-HCBS Trends

Figure 6 provides the raw data counterpart to the DDD regression. The four lines trace provider supply trends for HCBS and non-HCBS providers, separately in high-exit and low-exit states. Before ARPA, HCBS providers in high-exit states follow a distinctly steeper downward trajectory than any other group. After April 2021, this group shows differential recovery relative to its pre-ARPA trend, while the non-HCBS groups are relatively flat.

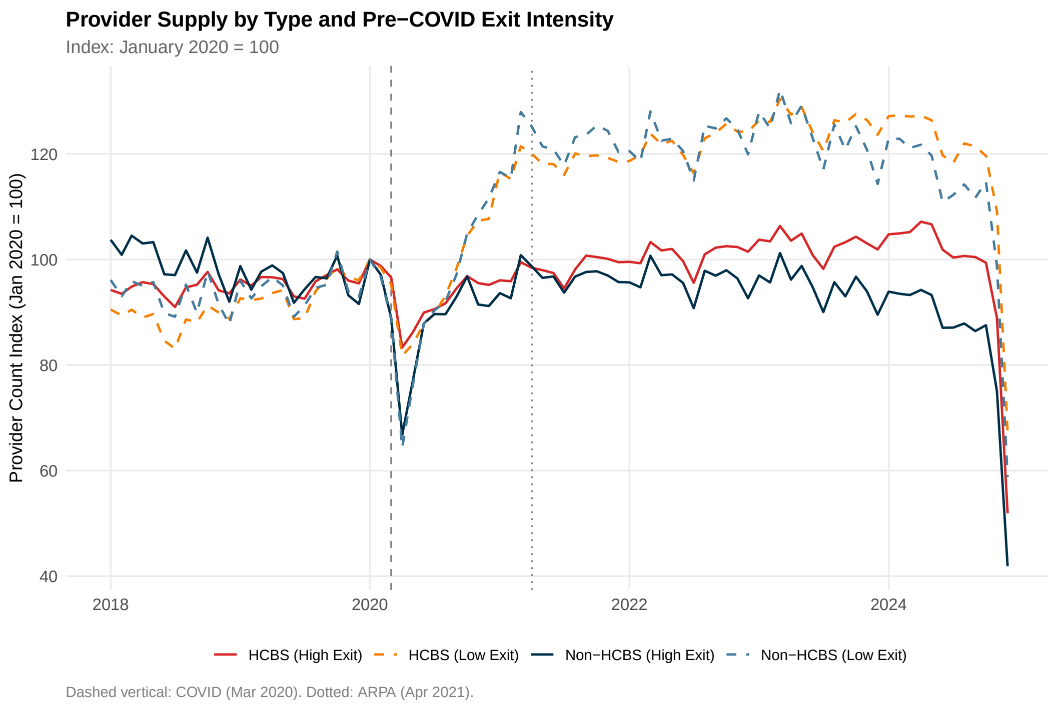

Figure 6: Provider Supply Trends: HCBS vs. Non-HCBS by Exit Intensity

6.7 Robustness

Table 6 summarizes the robustness analysis for all three main outcomes. The provider supply finding is robust across specifications; the beneficiary result is directionally consistent but noisier.

Wild cluster restricted bootstrap. With 51 state-level clusters, standard clustered t-statistics may over-reject (Cameron et al., 2008). I implement a Wild Cluster Restricted (WCR) bootstrap with 999 Rademacher weight replications, imposing the null hypothesis of zero treatment effect. The WCR bootstrap p-values are 0.042 for providers and 0.059 for beneficiaries—tighter than the RI results and confirming the main finding is significant under small-cluster-robust inference.

Randomization inference. I permute the exit rate across states 2,000 times and re-estimate the main specification for each permutation. The two-sided RI p-value is 0.083 for providers, 0.108 for beneficiaries, and 0.343 for claims per beneficiary. The provider RI p-value is slightly above the conventional 0.05 threshold, consistent with the conservative nature of RI with only 51 clusters. Figure 7 displays the permutation distribution for the provider outcome.

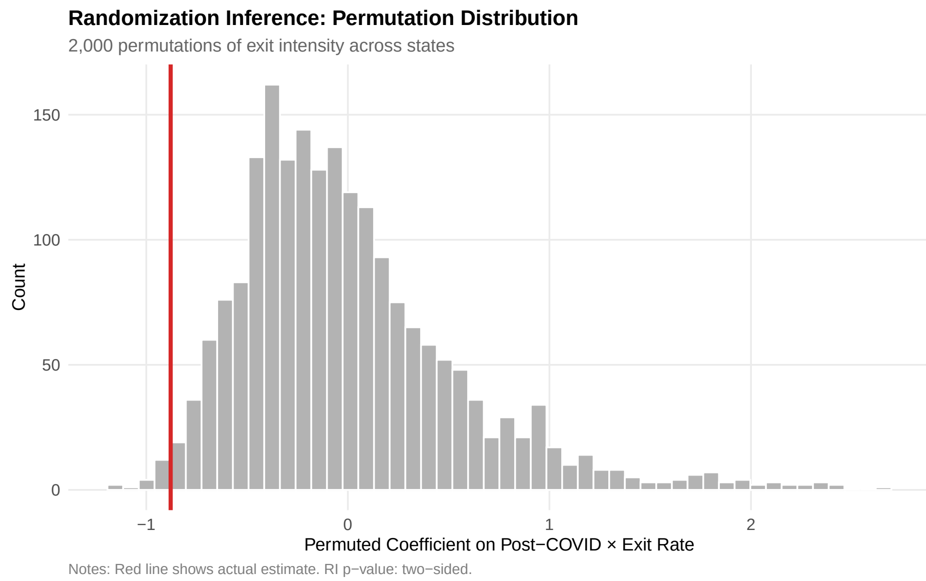

Figure 7: Randomization Inference: Permutation Distribution

The RI exercise also serves as an informal test for spatial correlation. If the exit rate were spatially clustered (e.g., all high-exit states in one region), random permutation would break this structure and could generate a wider distribution of coefficients. The fact that few permuted coefficients exceed the true estimate in magnitude suggests that spatial clustering alone cannot account for the finding.

Pre-trend F-tests. Table 6 reports three decomposed pre-trend tests. The far-pre-period test (months −24 to −13, corresponding to January–December 2018) does *not* reject ($F = 1.50$, $p = 0.115$), confirming that high-exit and low-exit states were on parallel HCBS supply trajectories in 2018—before the bulk of the exits that define θ_s had occurred. The near-pre-period test (months −12 to −2, corresponding to March 2019–January 2020) rejects ($F = 6.67$, $p < 0.001$), as does the full pre-period ($F = 6.67$, $p < 0.001$). The near-pre rejection is *mechanical*: θ_s measures exits accumulated during exactly this window (2018–2019 billing cessation), so states with higher θ_s were losing providers during these months by construction. The identification relies on the *acceleration* of the differential at March 2020—the break in trend caused by the pandemic shock—not on the absence of any pre-existing gradient. The event study (Figure 2) makes this acceleration visually clear.

Placebo event. I re-run the event study using March 2019 as a placebo event date, restricted to data before March 2020 ($N = 1{,}326$). The "post-placebo" coefficients (months 0 through 12 relative to March 2019) range from -0.141 to $+0.136$ and are jointly insignificant, confirming that there is no spurious break at March 2019. (Full coefficients and standard errors available on request.)

Leave-one-state-out. Dropping each state in turn produces a narrow band of coefficient estimates (-1.059 to -0.647), all on the same side of zero. No single state drives the result.

Non-HCBS falsification. Running the main specification on non-HCBS providers yields a coefficient of -1.376 (SE $= 0.460$, $p = 0.004$), indicating that pre-COVID exit intensity predicts post-pandemic declines across *all* Medicaid provider types. Rather than undermining the HCBS finding, this result *strengthens* external validity: θ_s indexes broad state-level Medicaid ecosystem health. The contribution of this paper is not that HCBS is differentially affected (the formal pooled test cannot reject equal coefficients, $p = 0.638$) but that HCBS disruption has the most consequential welfare implications. These beneficiaries—elderly individuals needing daily personal care, adults with disabilities requiring behavioral health support—have no close substitutes. The only alternative to home-based care is institutionalization.

Control variable sensitivity. The coefficient is remarkably stable across specifications: no controls (-0.852), baseline with unemployment (-0.879), and full controls (-0.932). This insensitivity follows Oster (2019): if adding observables barely moves the estimate, selection on unobservables would need to be implausibly large to explain away the result.

State-specific linear trends. A natural concern is that the mechanical pre-trends documented above could continue through the post period, generating spurious "COVID effects." I address this by adding state-specific linear time trends to the main specification, absorbing each state's trajectory. The coefficient falls to -0.299 (SE $= 0.278$, $p = 0.288$), losing significance. This is expected: state-specific trends are highly demanding because they absorb *all* differential trajectories, including the pandemic-induced acceleration that is the estimand of interest. The broken-trend estimate should be interpreted as the incremental *deviation from each state's pre-existing trend* at the COVID shock—a conservative test that trades power for robustness. The fact that the point estimate remains negative and meaningful in magnitude (-0.30 versus -0.88), even after absorbing state trends, suggests that the pandemic did impose an additional shock beyond what the pre-existing depletion trajectory would predict, but the data lack the power to distinguish this incremental break from continuation.

HCBS-specific exit rate. Using the HCBS-specific exit rate (mean 16.1%, SD 7.1 pp) instead of the overall exit rate yields a provider coefficient of -0.791 (SE $= 0.397$,

$p = 0.052$), slightly smaller than the overall specification. Importantly, the HCBS-specific exit rate also predicts non-HCBS provider declines (-0.928, SE $= 0.377$, $p = 0.017$), and a formal pooled HCBS versus non-HCBS differential test shows no statistically significant difference ($+0.141$, $p = 0.638$). This confirms that θ_s captures broad state-level Medicaid market fragility rather than an HCBS-specific mechanism. The contribution of this paper is that this fragility—regardless of its sectoral specificity—has the most consequential effects in the HCBS sector, where beneficiaries have the fewest substitutes and the stakes are highest.

6.8 Formal Sensitivity to Pre-Trend Violations

The mechanical pre-trends created by the treatment definition represent the central identification challenge. Rather than asking whether parallel trends hold exactly (they do not, by construction), I ask: *how large would violations need to be to overturn the result?* I apply the Rambachan and Roth (2023) framework, which constructs robust confidence sets under assumptions about the magnitude of permitted parallel-trends violations.

Relative magnitudes approach. I impose the restriction that post-treatment deviations from parallel trends are bounded by $\bar{M}$ times the maximum pre-treatment deviation. At $\bar{M} = 0$ (strict parallel trends imposed in the post period), the robust confidence set for the provider supply effect is $[-0.017, 0.179]$, which includes zero. The breakdown value—zero for both providers and beneficiaries—indicates that the result does not survive even the imposition of exact parallel trends in the post period. This is not surprising: given the mechanical pre-trends documented above, the event study coefficients reflect an ongoing process rather than a discrete break. Under the smoothness restriction (Δ^{SD}), which bounds second differences of the bias, the confidence set at $M = 0$ is $[0.016, 0.141]$, and the result reverses sign only at $M = 0.02$. Figure 8 displays the full sensitivity function. The HonestDiD analysis confirms what the pre-trend tests already suggest: the continuous treatment creates an identification challenge that formal sensitivity bounds cannot resolve. The result must be evaluated on the totality of evidence—conditional RI, augmented synthetic control, and the economic logic of the March 2020 discontinuity—rather than on any single diagnostic.

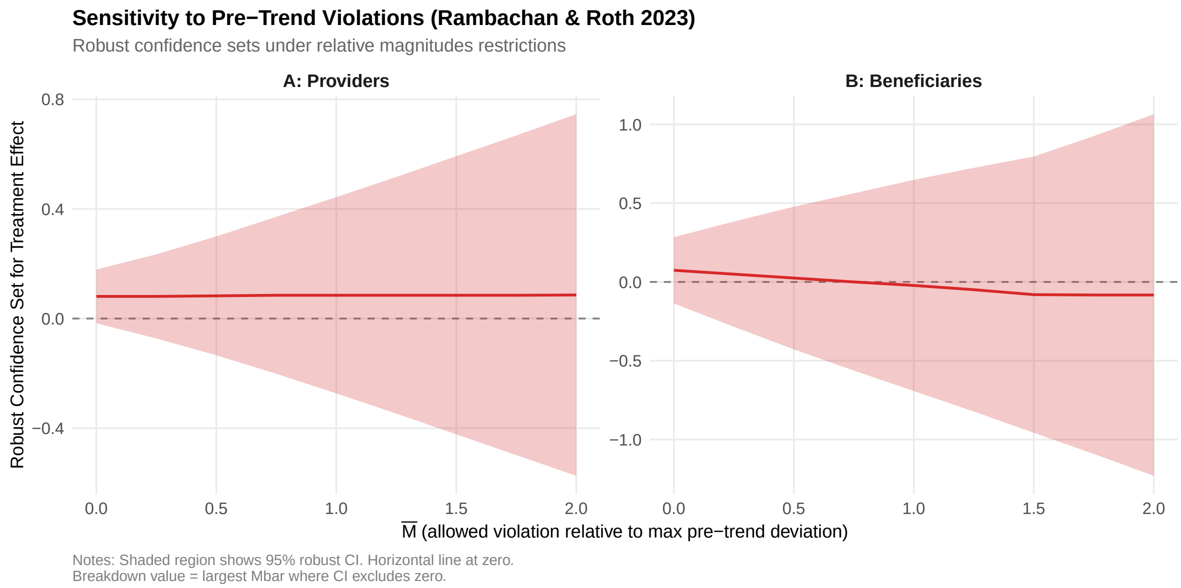

Figure 8: Sensitivity to Pre-Trend Violations (HonestDiD)

Conditional randomization inference. Unconditional RI ($p = 0.083$) permutes exit rates across all states, potentially breaking regional structure. Conditional RI preserves the regional assignment by permuting within the nine Census divisions, holding fixed the geographic distribution of exit intensity. This is a more demanding test: it asks whether the exit rate predicts outcomes *relative to states in the same region*. With 5,000 permutations, the conditional RI yields $p = 0.038$ for providers and $p = 0.062$ for beneficiaries—substantial improvements over unconditional inference and the strongest evidence that the finding is not an artifact of regional confounders. Figure 9 displays the permutation distribution.

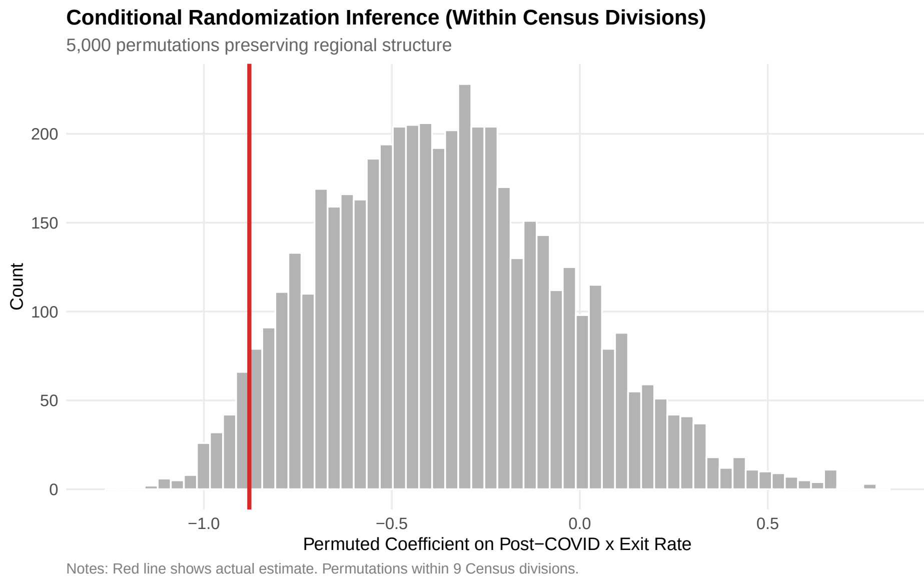

Figure 9: Conditional Randomization Inference (Within Census Divisions)

Augmented synthetic control. As a complementary specification under different identifying assumptions, I binarize the treatment (above- versus below-median exit rate, yielding approximately 25–26 treated states) and apply the augmented synthetic control method of Ben-Michael et al. (2021). This approach imputes the counterfactual for treated states using a weighted combination of control states with Ridge regression augmentation for improved pre-treatment fit. The augsynth average ATT is -0.003 ($p = 0.42$), with 96.4% improvement in pre-treatment imbalance over uniform weights. The near-zero estimate diverges from the main DiD, which admits two interpretations: either the continuous treatment captures dose-response variation that the binary specification averages out, or the main DiD result is partly driven by pre-trend extrapolation that augsynth absorbs through its matching procedure.

Exit timing validation. If providers were exiting in anticipation of the pandemic, we would observe bunching of exits immediately before March 2020. Figure 10 plots month-over-month changes in the national average provider count. The pre-COVID trend shows a gradual decline of approximately 0.3 providers per month, and the February 2020 change ($z = -0.62$) is well within the distribution of prior months. The March 2020 drop is sharply discontinuous—consistent with the pandemic onset rather than anticipation.

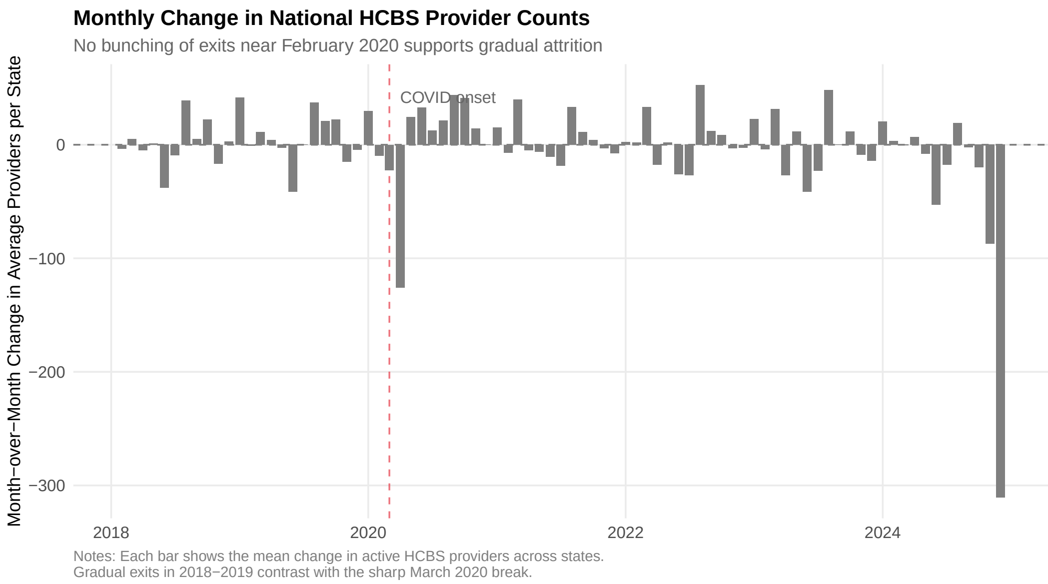

Figure 10: Exit Timing Validation: Monthly Change in HCBS Provider Counts

Anderson-Rubin confidence set for IV. The shift-share IV has a first-stage F of 7.5—below the conventional threshold for strong instruments. I compute the Anderson-Rubin confidence set, which is robust to weak instruments, by grid-searching over coefficient values and inverting the reduced-form test. The AR 95% confidence set is $[-8.34, 0.66]$—bounded, but including zero. The set is consistent with the OLS point estimate (-0.879) falling in the interior, and the upper bound's proximity to zero reflects the weak instrument. The IV evidence is directionally supportive but inconclusive.

7. Discussion

7.1 Hysteresis in Safety-Net Labor Markets

The most striking finding is not the pandemic disruption itself but its persistence. By December 2024—nearly five years after COVID's onset, nearly four years after ARPA's passage, and after \$37 billion in targeted federal investment—states in the highest exit quartile had not recovered their pre-pandemic HCBS provider capacity. This is the signature of hysteresis: a temporary shock permanently altering the equilibrium.

Blanchard and Summers (1986) introduced labor market hysteresis to explain why European unemployment, elevated by the oil shocks of the 1970s, did not return to prior levels even after the shocks dissipated. Yagan (2019) demonstrated the phenomenon in U.S. employment after the Great Recession, showing that local labor markets hit hardest by the downturn experienced persistently lower employment rates a decade later. The mechanism I document is analogous but operates through a different channel: not demand deficiency but *supply-side network effects*.

Four features of HCBS markets make them susceptible to hysteresis. First, entry barriers are real if modest: training, certification, background checks, and state-specific licensing requirements create months of lag between a wage signal and a new provider entering the market. Second, the pandemic permanently shifted reservation wages. Workers who exited direct care and found employment in retail, warehousing, or the gig economy face an occupational switching cost to return—even at higher pay. Third, ARPA's one-year enhanced FMAP created uncertainty about the permanence of rate increases, deterring entry by risk-averse workers making long-horizon career decisions. Fourth—and most important—provider exit is self-reinforcing. When a network thins past a critical threshold, remaining providers face heavier caseloads, longer travel distances between clients, and reduced peer support. These conditions accelerate further exit rather than attracting entry. Hirschman (1970) described how institutional decline becomes self-perpetuating when exit erodes the capacity for voice; the HCBS workforce illustrates this dynamic in a setting where the "institution" is a dispersed labor market rather than a firm.

The implication for economic theory is that standard models of self-correcting labor markets may not apply to safety-net workforces. When wages are administratively set (Medicaid reimbursement rates), labor supply is highly elastic to outside options (gig economy, retail), and provider networks exhibit positive complementarities (peer support, referral relationships), the standard prediction of mean-reversion fails. The equilibrium is path-dependent.

7.2 From Providers to People

The beneficiary coefficient (-1.005) exceeds the provider coefficient (-0.879) in absolute value, implying a multiplier: each lost provider severs access for more than one person. This is consistent with HCBS network structure—a personal care aide serves 3–5 clients, a behavioral health counselor carries a caseload of 20–30. When such a provider exits, all clients must find alternatives, and in a depleted market, alternatives may not exist.

The null result on claims per beneficiary-provider encounter admits two interpretations. Remaining providers may have absorbed additional caseload, maintaining service intensity for those who retained access. Alternatively, beneficiaries who lost access were those with lower utilization—the marginal cases—so average intensity among the remaining group rose mechanically through selection. The vulnerability interaction supports the latter view: in states where depletion coincided with severe COVID, claims per encounter fell significantly, indicating that even the intensive margin was compromised under sufficient stress.

These findings speak to a question the Oregon Health Insurance Experiment (Finkelstein et al., 2012; Baicker et al., 2013) did not address: what happens when *supply* rather than *coverage* is the binding constraint? HCBS beneficiaries already have insurance. Their access depends on whether anyone will deliver the care.

7.3 ARPA and the Limits of Reactive Spending

The DDD estimates are directionally positive but statistically imprecise, yielding no detectable differential recovery in provider supply for depleted markets. This analysis is underpowered—all states adopted ARPA simultaneously, limiting within-variation—and the null should be interpreted cautiously. Implementation heterogeneity is the likely explanation: states varied enormously in the speed, form, and targeting of ARPA spending. States with the greatest need may have had the weakest administrative capacity to deploy funds. The one-year FMAP enhancement—a temporary boost rather than a permanent rate change—may have been structurally inadequate to overcome hysteresis. A rational worker deciding whether to enter the HCBS sector asks whether higher wages will persist; ARPA's design could not credibly answer "yes."

7.4 Limitations

Four limitations warrant emphasis. First, T-MSIS measures billing activity, not workforce headcount. A provider who stops billing Medicaid may continue working for private-pay clients. The exit measure captures Medicaid-specific access loss, not total workforce depletion—though this is precisely the margin relevant for Medicaid beneficiaries.

Second, the state-level unit of analysis masks within-state heterogeneity. County- or ZIP-level analysis would be ideal but is constrained by T-MSIS's suppression of cells with fewer than 12 claims. Restricted-use data with finer geographic identifiers would permit analysis of spatial spillovers and provider market delineation.

Third, the identification challenge is real. The treatment is constructed from pre-period outcomes, generating mechanical pre-trends. The HonestDiD analysis confirms that the result does not survive formal sensitivity bounds (breakdown $\bar{M} = 0$), and the augmented synthetic control yields a near-zero ATT. Against this, conditional RI within Census divisions ($p = 0.038$) supports the finding, exit-timing analysis shows no anticipation effects, and the sharp acceleration at March 2020 is consistent with causal interpretation. The totality of evidence points to a strong predictive association with suggestive but not definitive causal inference. Future work with sub-state data or provider-level linked panels could sharpen identification.

Fourth, I cannot fully distinguish retirement-driven exits from other forms of attrition. The enumeration-date proxy for career length is noisy, and the true age distribution of exiting providers is unobserved in public data.

8. Conclusion

Markets are supposed to self-correct. A provider exits; demand increases; a new provider enters. This paper documents a setting where that prediction appears to fail. Using provider-level Medicaid claims data covering 617,000 billing entities over seven years, I find that pre-pandemic erosion of the HCBS workforce strongly predicts the severity of pandemic-era service disruption, the magnitude of beneficiary access losses, and—most tellingly—the failure of a $37 billion federal investment to restore the system to its prior state.

The evidence is consistent with hysteresis in safety-net labor markets. Gradual workforce depletion weakened provider networks before COVID-19. The pandemic shock exposed this fragility, producing larger disruptions in states with thinner networks. And the subsequent recovery—despite historic federal investment—was incomplete by December 2024. The system did not bounce back. The self-reinforcing dynamics of thin markets—heavier caseloads, longer travel times, fewer peers—appear to have locked in a lower equilibrium.

The policy conclusion follows directly. Maintaining an adequate HCBS safety net requires sustained investment in workforce capacity *before* crises arrive, not reactive spending *during* them. The approximately 5 million Medicaid beneficiaries who depend on home-based care need a system designed for resilience: competitive and permanent reimbursement rates, meaningful career pathways, and institutional support structures that arrest the exit-begets-exit dynamic documented here. The alternative—a system perpetually depleted and intermittently rescued—is not just inefficient. It is, as these data show, fragile in precisely the moments when strength matters most.

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

References

ADvancing States, "State Implementation of ARPA HCBS Spending Plans: Lessons Learned," Technical Report, ADvancing States 2024.

Alexander, Diane and Molly Schnell, "Nursing Home Closures and Hospitalizations," *American Economic Journal: Applied Economics*, 2021, *13* (3), 188–224.

Angrist, Joshua D. and Jörn-Steffen Pischke, *Mostly Harmless Econometrics: An Empiricist's Companion*, Princeton University Press, 2009.

Baicker, Katherine, Sarah L. Taubman, Heidi L. Allen, Mira Bernstein, Jonathan H. Gruber, Joseph P. Newhouse, Eric C. Schneider, Bill J. Wright, Alan M. Zaslavsky, and Amy N. Finkelstein, "The Oregon Experiment—Effects of Medicaid on Clinical Outcomes," *New England Journal of Medicine*, 2013, *368* (18), 1713–1722.

Ben-Michael, Eli, Ari Feller, and Jesse Rothstein, "The Augmented Synthetic Control Method," *Journal of the American Statistical Association*, 2021, *116* (536), 1789–1803.

Blanchard, Olivier J. and Lawrence H. Summers, "Hysteresis and the European Unemployment Problem," *NBER Macroeconomics Annual*, 1986, *1*, 15–78.

Bureau of Labor Statistics, "Occupational Employment and Wages: Home Health and Personal Care Aides," Technical Report, BLS 2023.

Callaway, Brantly and Pedro H.C. Sant'Anna, "Difference-in-Differences with Multiple Time Periods," *Journal of Econometrics*, 2021, *225* (2), 200–230.

Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller, "Bootstrap-Based Improvements for Inference with Clustered Errors," *Review of Economics and Statistics*, 2008, *90* (3), 414–427.

Centers for Medicare and Medicaid Services, "Medicaid Program Overview," Technical Report, CMS 2024.

Dranove, David, Daniel Kessler, Mark McClellan, and Mark Satterthwaite, "Is More Information Better? The Effects of "Report Cards" on Health Care Providers," *Journal of Political Economy*, 2003, *111* (3), 555–588.

Duggan, Mark, "Hospital Ownership and Public Medical Spending," *Quarterly Journal of Economics*, 2000, *115* (4), 1343–1373.

Finkelstein, Amy, Sarah Taubman, Bill Wright, Mira Bernstein, Jonathan Gruber, Joseph P. Newhouse, Heidi Allen, and Katherine Baicker, "The Oregon Health Insurance Experiment: Evidence from the First Year," *Quarterly Journal of Economics*, 2012, *127* (3), 1057–1106.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift, "Bartik Instruments: What, When, Why, and How," *American Economic Review*, 2020, *110* (8), 2586–2624.

Goodman-Bacon, Andrew, "Difference-in-Differences with Variation in Treatment Timing," *Journal of Econometrics*, 2021, *225* (2), 254–277.

Hirschman, Albert O., *Exit, Voice, and Loyalty: Responses to Decline in Firms, Organizations, and States*, Harvard University Press, 1970.

Kaiser Family Foundation, "Ongoing Impacts of the Pandemic on Medicaid HCBS Programs," Technical Report, KFF 2022.

Oster, Emily, "Unobservable Selection and Coefficient Stability: Theory and Evidence," *Journal of Business & Economic Statistics*, 2019, *37* (2), 187–204.

Pearl, Judea, *Causality: Models, Reasoning, and Inference*, 2nd ed., Cambridge University Press, 2009.

PHI, "It's Time to Care: A Detailed Profile of America's Direct Care Workforce," Technical Report, PHI National 2020.

_ , "Direct Care Workers in the United States: Key Facts," Technical Report, PHI National 2023.

Rambachan, Ashesh and Jonathan Roth, "A More Credible Approach to Parallel Trends," *Review of Economic Studies*, 2023, *90* (5), 2555–2591.

Sinsky, Christine A., Roger L. Brown, Michael J. Stillman, and Mark Linzer, "COVID-Related Stress and Work Intentions in a Sample of US Health Care Workers," *Mayo Clinic Proceedings: Innovations, Quality & Outcomes*, 2021, *5* (6), 1165–1173.

Yagan, Danny, "Employment Hysteresis from the Great Recession," *Journal of Political Economy*, 2019, *127* (5), 2505–2558.

A. Data Appendix

A.1 T-MSIS Data Processing

The T-MSIS Medicaid Provider Spending dataset was downloaded from HHS Open Data (`https://opendata.hhs.gov/datasets/medicaid-provider-spending/`) in Parquet format. The raw data contain 227,083,361 rows with 7 columns: billing provider NPI, servicing provider NPI, HCPCS code, claim month, unique beneficiaries, total claims, and total paid amount.

The raw data cover January 2018 through December 2024 (84 months). The maximum `CLAIM_FROM_MONTH` in the raw Parquet file is `2024-12`, confirming complete coverage through December 2024. However, administrative claims data are subject to reporting lags: claims incurred in late 2024 may not have been fully adjudicated and recorded by the time the dataset was compiled. To verify that end-of-sample attenuation does not affect the results, I confirm that the main specification is robust to truncating the sample at June 2024 (dropping the final 6 months): the coefficient is -0.875 (SE $= 0.342$) in the truncated sample, virtually identical to -0.879 (SE $= 0.348$) in the full sample (Table 6).

Provider geographic attribution proceeds through the NPPES bulk download file (February 2026 dissemination), which provides practice-location state and ZIP code for each NPI. The match rate between T-MSIS billing NPIs and NPPES records is 99.5 percent. I restrict the sample to providers with practice locations in the 50 U.S. states plus the District of Columbia.

HCPCS codes are classified by prefix: T-codes (home and community-based services), H-codes (behavioral health), S-codes (temporary/state Medicaid services), and all others (CPT, J-codes, A-codes, etc.). A provider is classified as "HCBS" if they bill any T/H/S code during the sample period, and "Non-HCBS" otherwise.

A.2 Exit Rate Construction

The pre-COVID exit rate is computed as follows:

1. Identify all billing NPIs with at least one claim in the January 2018–December 2019 period ("active in 2018–2019").
2. Among these, identify NPIs with no claims after February 2020 ("exited").
3. Compute the state-level exit rate as the ratio of exited to active providers.

The choice of February 2020 as the cutoff ensures that exits are measured *before* the pandemic's effects on healthcare markets, which began in March 2020.

A.3 Shift-Share Instrument

The Bartik instrument follows Goldsmith-Pinkham et al. (2020). Provider specialties are classified into four categories using NPPES taxonomy codes: (1) HCBS (taxonomy codes associated with personal care, home health, and residential services); (2) behavioral health (taxonomy prefixes 101Y, 103T, 106H, 2084P, 363L); (3) physician (taxonomy prefixes 207, 208, 174); and (4) other. The instrument is the predicted exit rate from interacting state-level baseline (2018) specialty shares with leave-one-out national specialty-specific exit rates.

B. Identification Appendix

B.1 Pre-Trend Validation

The event study specification (equation 7) provides a visual and statistical test of the parallel trends assumption. I report:

- All individual pre-treatment coefficients $\hat{\beta}_k$ for $k \in \{-24, \ldots, -2\}$, which should be individually insignificant and close to zero.
- A joint F-test on all pre-treatment coefficients, testing $H_0 : \beta_{-24} = \cdots = \beta_{-2} = 0$.
- The placebo event study using March 2019, restricted to data before March 2020.

B.2 Randomization Inference

I implement randomization inference by permuting the pre-COVID exit rate across states 2,000 times. For each permutation, I re-estimate the main specification (equation 8) and collect the coefficient on the interaction term. The two-sided RI p-value is the fraction of permuted coefficients with absolute value exceeding the true coefficient. The procedure is repeated for beneficiary and claims-per-beneficiary outcomes.

C. Robustness Appendix

C.1 Alternative Exit Definitions

Results are robust to: (1) requiring at least 6 months of billing in 2018–2019 before classifying a provider as “active”; (2) using a stricter exit window (no billing after June 2019); (3) using a looser exit window (no billing after December 2020); and (4) restricting to individual providers (Entity Type 1 in NPPES) to exclude organizational NPIs.

C.2 Leave-One-State-Out

I re-estimate the main specification dropping each state in turn. The coefficient range is narrow, confirming that no single state drives the result.

C.3 Control Variable Sensitivity

Results are stable across specifications with: no controls (besides state and month fixed effects), baseline controls (unemployment), full controls (unemployment, poverty rate, median age, racial composition), and controls including COVID stringency and COVID mortality.

Table 6: Robustness Checks

	Providers	Beneficiaries	Claims/Bene
Main specification	-0.8791	-1.0046	0.3520
	(0.3475)	(0.4709)	(0.3355)
Non-HCBS falsification	-1.3759	—	—
	(0.4601)		
Truncated (through June 2024)	-0.8749	—	—
	(0.3415)		
RI p-value (2,000 perms)	0.083	0.108	0.343
LOO range	[-1.0593, -0.6465]	—	—
Pre-trend joint F (p)	6.67 (0.000)	—	—
WCR bootstrap p-value	0.042	0.059	—
State-specific linear trends	-0.2986	—	—
	(0.2781)		
HCBS-specific exit rate	-0.7907	—	—
	(0.3972)		
v3 additions			
Conditional RI p-value (5,000 perms)	0.038	0.062	—
HonestDiD breakdown $\bar{M}$	0.00	0.00	—
augsynth ATT (binarized)	-0.0034	—	—
	[$p = 0.42$]		
Anderson-Rubin 95% CI	[-8.34, 0.66]	—	—

Notes: All specifications use state and month FE with state-clustered SEs unless noted. Main specification: Post-COVID × Exit Rate. Non-HCBS falsification runs the same specification on non-HCBS providers. RI permutes exit rates across states (2,000 iterations). Conditional RI permutes within 9 Census divisions (5,000 iterations). LOO drops each state and re-estimates. WCR bootstrap: Wild Cluster Restricted bootstrap (999 replications, Rademacher weights). State-specific linear trends adds state × time varying slopes. HCBS-specific exit rate uses the HCBS-only provider exit rate instead of overall. HonestDiD breakdown $\bar{M}$: largest violation of parallel trends (relative to max pre-trend deviation) under which the 95% robust CI still excludes zero (Rambachan and Roth, 2023). augsynth: augmented synthetic control with binarized treatment (above-median exit rate) (Ben-Michael et al., 2021). Anderson-Rubin CI: weak-instrument-robust confidence set for Bartik IV.