

Where Medicaid Goes Dark: A Claims-Based Atlas of Provider Deserts and the Resilience of Supply to Enrollment Shocks

APEP Autonomous Research* @SocialCatalystLab

February 20, 2026

Abstract

Using newly released T-MSIS Medicaid claims data covering 227 million provider-service records (2018Q1–2024Q3), we construct the first claims-based atlas of county-level Medicaid provider deserts across six clinical specialties. A key methodological contribution is an all-clinicians measure that maps nurse practitioners and physician assistants to their clinical specialties—reflecting that NPs deliver the majority of Medicaid primary care in rural areas. Even with NPs included, the descriptive portrait is alarming: over 99% of county-quarters fall below the desert threshold for psychiatry, 89% are classified as primary care deserts, and rural desert rates remain extreme. Exploiting staggered state-level variation in post-pandemic enrollment unwinding, we find that provider supply is *remarkably inelastic* to short-run demand shocks: estimates are precisely estimated near zero across all specialties, confirmed by permutation inference. Medicaid deserts reflect chronic structural factors—plausibly including reimbursement and workforce dynamics—rather than acute enrollment fluctuations.

JEL Codes: I11, I13, I18, J44

Keywords: Medicaid, provider access, medical deserts, enrollment unwinding, health workforce, supply inelasticity, nurse practitioners

*Autonomous Policy Evaluation Project. This paper is a revision of apep_0417 v2 (https://github.com/SocialCatalystLab/ape-papers/tree/main/apep_0417). Total execution time: 1h 34m (cumulative: 1h 34m). Correspondence: scl@econ.uzh.ch

1. Introduction

In April 2023, American states began the largest administrative shock to health insurance in the nation’s history. The end of the pandemic-era continuous enrollment provision triggered the “unwinding” of Medicaid rolls, eventually removing more than 25 million people from coverage (Corallo and Moreno, 2024). The policy debate has understandably focused on those who lost insurance. But a natural concern extends to the supply side: when patients vanish, do their doctors follow? If enrollment loss drives provider exit, the unwinding could deepen medical deserts, trapping those who remain covered in areas where care is increasingly unavailable.

This paper investigates this question using the most comprehensive data ever assembled on Medicaid provider activity, and arrives at a surprising answer. The Medicaid provider crisis is real — and severe. But the enrollment unwinding is not its cause.

The question matters because Medicaid access has always been about more than an insurance card. Decades of research have documented a persistent gap between formal eligibility and realized access. Decker (2012) showed that only two-thirds of physicians accept new Medicaid patients, a rate that has declined over time. Polsky et al. (2015) found that simulated Medicaid patients seeking appointments were turned away at rates far exceeding the privately insured. The fundamental driver is reimbursement: Medicaid pays physicians roughly 72 cents on the Medicare dollar, creating a financial disincentive that concentrates willing providers in densely populated areas where volume compensates for low margins (Zuckerman et al., 2009). The result is that Medicaid beneficiaries in rural and low-income communities face severe shortages that existing measurement tools systematically understate.

The standard measure of physician shortage — the Health Resources and Services Administration’s Health Professional Shortage Area (HPSA) designation — relies on provider-to-population ratios derived from survey data and registries (Health Resources and Services Administration, 2023). These data capture who *could* treat Medicaid patients, not who actually does. A county may have adequate physician supply for its privately insured residents while functioning as a complete desert for Medicaid beneficiaries. The distinction between registered providers and active billers is empirically enormous and, until now, unmeasurable at scale.

This paper exploits the recent release (February 2026) of the Transformed Medicaid Statistical Information System (T-MSIS) provider spending data to make two contributions. The first is descriptive: we construct the first claims-based atlas of Medicaid provider deserts. T-MSIS records every provider-service transaction processed through state Medicaid programs, yielding 227 million billing provider–servicing provider–procedure code observations spanning

2018Q1–2024Q3. We link these claims to the National Plan and Provider Enumeration System (NPES) to classify providers by specialty and location, constructing a county \times specialty \times quarter panel covering 3,144 continental US counties and six clinical specialty groups: primary care, psychiatry, behavioral health, dental, obstetrics-gynecology, and surgery.

A key methodological innovation is our *all-clinicians* provider measure, which maps nurse practitioners (NPs) and physician assistants (PAs) to their clinical specialties using NUCC taxonomy subcategories. Prior work has typically counted NPs as a separate category or excluded them entirely. This understates the effective Medicaid provider supply in specialties where NPs deliver a substantial share of care. Family NPs, the single largest NP subcategory, are classified under primary care; psychiatric-mental health NPs under psychiatry; women’s health NPs under OB-GYN. We retain a no-NP/PA classification as a robustness measure, which excludes nurse practitioners and physician assistants but retains non-physician clinicians (psychologists, social workers, dentists) whose specialty classifications are unaffected by the NP/PA inclusion decision. The comparison between measures is informative: including NPs reduces primary care desert rates substantially while leaving psychiatry and OB-GYN deserts nearly unchanged.

The descriptive facts are striking. Even with the all-clinicians measure, 99.6% of county-quarters fall below our desert threshold for psychiatry (fewer than one active clinician per 10,000 population). Approximately 89% are classified as primary care deserts, and 99.8% as OB-GYN deserts. These are not registry counts of who holds a Medicaid provider number — they reflect who actually submitted claims. The declines are not uniform: rural counties experience desert rates two to three times higher than urban counties across every specialty. By 2023, vast contiguous regions of rural America had zero active Medicaid coverage for specialty care. The comparison between our all-clinicians and no-NP/PA measures reveals that NP inclusion substantially reduces primary care deserts but has minimal impact on psychiatry, OB-GYN, and surgery deserts — specialties where NP substitution remains limited by credentialing and scope-of-practice constraints.

The second contribution is causal — and its main finding is a null. We exploit the staggered state-level variation in Medicaid enrollment unwinding to estimate whether the demand shock of disenrollment drives provider exit. States varied dramatically in both timing and intensity: net enrollment losses ranged from 1.4% in Maine to 30.2% in Colorado, with a median of 14.0% ([Kaiser Family Foundation, 2024](#)). Our difference-in-differences design estimates the effect of a state’s cumulative disenrollment rate on county-level active provider counts, validated by event-study evidence of flat pre-trends across eight pre-treatment quarters.

The central result is that provider supply is remarkably inelastic to the enrollment

shock. The pooled estimate across all specialties is precisely estimated near zero. No individual specialty shows a statistically significant response. The event study confirms clean pre-trends and shows no meaningful post-treatment departure from zero. Permutation inference with 500 randomizations of state treatment timing places our observed coefficient squarely in the middle of the null distribution. The null result is robust to alternative active-provider thresholds, a full-time provider definition (≥ 36 claims per quarter), binary treatment definitions, controlling for time-varying Medicaid share, region \times quarter fixed effects, a Medicaid-population denominator, total claims as the outcome variable, and a pre-period placebo test. The null holds identically when estimated using the no-NP/PA panel, confirming it is not an artifact of the NP inclusion decision.

The null finding is important precisely because it is surprising. The intuitive hypothesis — that losing 25 million patients should cause providers to exit — has a clear theoretical motivation and has shaped policy discussions around the unwinding. Our evidence suggests this mechanism does not operate on the margins observed in the data. This inelasticity is consistent with [Clemens and Gottlieb \(2014\)](#), who find that physician behavior responds primarily to reimbursement prices rather than volume changes, and with the general equilibrium insights of [Finkelstein \(2007\)](#) on insurance market expansions. Several additional explanations are consistent with the null: Medicaid typically constitutes a minority of a provider’s payer mix, provider participation decisions may respond to long-run structural factors rather than short-run enrollment fluctuations, and the unwinding may have been too recent for supply-side adjustments to fully materialize.

This paper contributes to several literatures. First, we advance the measurement of health care access by constructing the first claims-based provider desert atlas, complementing survey-based approaches ([Goodman et al., 2023](#); [Ricketts and Randolph, 2007](#); [Graves, 2016](#)). Our all-clinicians measure is the first to systematically incorporate NPs into specialty-specific desert designations using administrative claims, revealing both the importance of NP inclusion for primary care measurement and its limitations for specialty care. Second, we contribute to the literature on Medicaid provider participation, which has documented the roles of reimbursement rates ([Decker, 2012](#); [Zuckerman et al., 2009](#); [Gruber, 2003](#)), physician attitudes ([Long, 2013](#)), and managed care penetration ([Duggan, 2004](#)). Our null causal result suggests that the demand side — at least in the form of short-run enrollment shocks — is *not* a first-order determinant of provider participation. Third, we contribute to the rapidly growing literature on the Medicaid enrollment unwinding, which has documented coverage losses ([Corallo and Moreno, 2024](#); [Kaiser Family Foundation, 2024](#); [Wagner and Erzouki, 2023](#)) and early evidence on health care utilization ([Sommers et al., 2024](#)). We provide the first evidence on supply-side effects.

Our work is also related to the broader literature on public insurance and health care supply. [Garthwaite et al. \(2014\)](#) showed that public insurance crowd-out operates partly through the provider market. [Sommers et al. \(2017\)](#) documented that Medicaid expansion improved access to care. [Miller and Wherry \(2019\)](#) showed that Medicaid coverage reduces mortality, an effect that depends on provider availability. The maintained hypothesis in much of this literature — that provider supply responds elastically to demand — receives only limited support in our setting.

2. Background

2.1 The Medicaid Provider Access Problem

Medicaid is the largest source of health insurance in the United States, covering approximately 90 million people before the pandemic and peaking near 94 million during the continuous enrollment period ([Centers for Medicare and Medicaid Services, 2024](#)). Despite this scale, access to care for Medicaid beneficiaries has been a persistent policy concern. The core issue is reimbursement. Medicaid physician fees averaged 72% of Medicare rates nationally as of 2019, with wide interstate variation: New York paid 29% of Medicare for primary care office visits, while Alaska paid 175% ([Zuckerman et al., 2009](#); [Decker, 2023](#)). Low reimbursement discourages physician participation, creating a wedge between formal coverage and realized access.

The consequences are well documented. [Polsky et al. \(2015\)](#) conducted audit studies in which simulated patients called physicians seeking new-patient appointments. Medicaid callers were offered appointments in 58% of calls, compared with 85% for privately insured callers. [Decker \(2012\)](#) found that only 69% of physicians accepted new Medicaid patients nationally, with rates below 50% in some states. [Rhodes et al. \(2014\)](#) showed that areas with fewer Medicaid-accepting providers had worse access to preventive care and higher rates of emergency department utilization. The provider access problem is not merely an inconvenience; it determines whether insurance coverage translates into health care delivery.

2.2 Measuring Medical Deserts

The standard federal measure of health care access is the Health Professional Shortage Area (HPSA) designation, administered by HRSA since 1978. HPSA scores use provider-to-population ratios, poverty rates, and distance measures to classify geographic areas, populations, or facilities as underserved ([Health Resources and Services Administration, 2023](#)). The methodology has been criticized on several grounds: it relies on state-submitted

data that may be outdated, it counts all licensed physicians regardless of payer mix, and the geographic units (often counties or census tracts) may not reflect actual patient travel patterns ([Ricketts and Randolph, 2007](#); [Graves, 2016](#)).

For Medicaid-specific access, HPSA designations are particularly misleading. A county may have an adequate overall physician-to-population ratio while functioning as a complete desert for Medicaid patients if most physicians decline Medicaid. The distinction matters enormously for policy: HPSA scores drive billions in federal funding through programs like the National Health Service Corps, FQHC reimbursement bonuses, and J-1 visa waivers for international medical graduates ([Health Resources and Services Administration, 2023](#)). If these designations miss Medicaid-specific shortages, resources may be misdirected.

2.3 The Pandemic and Continuous Enrollment

The Families First Coronavirus Response Act (FFCRA) of March 2020 required states to maintain Medicaid enrollment for all beneficiaries as a condition of receiving enhanced federal matching funds. This “continuous enrollment” provision effectively suspended routine eligibility redeterminations, preventing disenrollment for any reason other than a beneficiary’s request, death, or interstate move. The result was a dramatic expansion of the Medicaid rolls. National enrollment grew from 71 million in February 2020 to an all-time high of approximately 94 million by March 2023, a 32% increase ([Centers for Medicare and Medicaid Services, 2024](#)).

The coverage expansion brought new patients into the Medicaid system, potentially supporting provider participation by increasing the volume of Medicaid patients in areas where provider networks had been thinning. At the same time, the pandemic itself disrupted health care delivery — elective procedures were postponed, telehealth expanded rapidly, and many providers experienced financial strain ([American Hospital Association, 2021](#)). The net effect on Medicaid provider supply during the pandemic era is ambiguous and represents part of the variation we document in our descriptive atlas.

2.4 The Unwinding

The Consolidated Appropriations Act of December 2022 set March 31, 2023, as the end date for the continuous enrollment provision. Beginning April 1, 2023, states could resume normal eligibility redeterminations and disenroll individuals found ineligible. The “unwinding” that followed was the largest administrative event in Medicaid’s history.

States varied dramatically in their approach. Early movers like Arkansas, Idaho, and South Dakota began processing renewals and issuing disenrollments immediately in April

2023. Oregon delayed until October 2023. California and New York processed renewals but moved cautiously, investing in outreach to minimize procedural disenrollments. By mid-2024, states had collectively removed over 25 million people from Medicaid rolls ([Corallo and Moreno, 2024](#); [Kaiser Family Foundation, 2024](#)).

The variation was not only in timing but in intensity. Net enrollment declines (from peak to trough) ranged from approximately 1.4% in Maine to 30.2% in Colorado, with a median state losing about 14% of peak enrollment. Critically, a large share of disenrollments were procedural — beneficiaries removed for failing to return paperwork rather than for confirmed ineligibility ([Wagner and Erzouki, 2023](#)). Several states paused their unwinding after discovering high procedural disenrollment rates, further adding to the cross-state variation in timing and intensity that we exploit.

2.5 Provider Market Dynamics and Nurse Practitioners

The Medicaid unwinding occurred against a backdrop of secular trends in the health care workforce. Physician retirement has accelerated, with the Association of American Medical Colleges projecting a shortage of up to 124,000 physicians by 2034 ([Association of American Medical Colleges, 2021](#)). The shortage is most acute in primary care and psychiatry, the specialties where Medicaid payment gaps relative to private insurance are largest. Simultaneously, the geographic concentration of physicians has intensified: [Goodman et al. \(2023\)](#) documented that rural areas lost physicians at a faster rate than urban areas throughout the 2010s, a trend driven by lifestyle preferences, institutional agglomeration effects, and the economics of low-volume practice.

Partially offsetting physician losses, nurse practitioners and physician assistants have expanded their scope of practice in many states. As of 2024, 27 states and the District of Columbia grant NPs full practice authority, allowing independent practice without physician oversight ([Alexander and Schnell, 2019](#); [American Association of Nurse Practitioners, 2024](#)). NPs are disproportionately likely to practice in rural and underserved areas, making them a critical component of the Medicaid safety net. In many rural counties, NPs are the *only* clinicians billing Medicaid for primary care. This fact motivates our all-clinicians provider measure: any assessment of Medicaid access that excludes NPs systematically overstates the extent of deserts, particularly in primary care and behavioral health where NP practice is most prevalent. Conversely, for specialties like psychiatry and OB-GYN, the number of NPs with relevant subspecialty credentials remains small relative to need, and NP inclusion has modest effects on measured desert rates.

3. Data

3.1 T-MSIS Provider Spending Data

Our primary data source is the Transformed Medicaid Statistical Information System (T-MSIS) provider spending file, released by the Department of Health and Human Services in February 2026. T-MSIS records every provider-service transaction processed through state Medicaid fee-for-service and managed care encounter systems. The raw data contain 227 million billing-NPI \times servicing-NPI \times HCPCS-code \times month observations spanning January 2018 through December 2024, though we restrict the analysis sample to 2018Q1–2024Q3 due to incomplete Q4 2024 data from billing lags. For each record, we observe the number of claims, total Medicaid payments, and the number of unique beneficiaries served.

We use the *billing* NPI as our unit of analysis for aggregation to the provider-quarter level. We aggregate from monthly to quarterly frequency to reduce noise from billing cycles and seasonal variation.

3.2 NPPES Provider Registry and Specialty Classification

We link T-MSIS records to the National Plan and Provider Enumeration System (NPPES), a CMS-maintained registry of all NPI holders. NPPES provides each provider’s practice location (mailing and practice-location ZIP codes), taxonomy codes (which encode specialty and credential type), and enumeration date.

We employ a *dual classification* approach using the National Uniform Claim Committee (NUCC) Health Care Provider Taxonomy codes. Our primary measure is an **all-clinicians** classification that maps nurse practitioners and physician assistants to their clinical specialties based on NUCC subcategory codes, reflecting the actual services these providers deliver. Our secondary measure is a **no-NP/PA** classification that restricts attention to physician specialties, used as a robustness check and to quantify the impact of NP inclusion on measured access.

The all-clinicians classification assigns providers to six specialty groups:

1. *Primary Care*: Family medicine, internal medicine, general practice, and pediatrics (MD/DO), plus family NPs, primary care NPs, adult health NPs, community health NPs, pediatric NPs, gerontology NPs, neonatal NPs, acute care NPs, and physician assistants (general and medical).
2. *Psychiatry*: Psychiatry and child/adolescent psychiatry (MD/DO), plus psychiatric-mental health NPs.

3. *Behavioral Health*: Psychologists, clinical social workers, licensed counselors, marriage and family therapists (credential-based, not affected by NP classification).
4. *Dental*: General dentistry, oral surgery, pediatric dentistry, orthodontics (credential-based).
5. *Obstetrics-Gynecology*: OB-GYN and maternal-fetal medicine (MD/DO), plus women’s health NPs.
6. *Surgery*: General, orthopedic, cardiovascular, neurological, plastic, colon/rectal, and thoracic surgery (MD/DO), plus surgical PAs.

This classification reflects a substantive judgment: NPs and PAs practicing under specialty-specific credentials are delivering the same type of care as their physician counterparts and should be counted as part of the available provider supply for that specialty. The specific NUCC taxonomy code mappings are detailed in Appendix Table 6.

3.3 Nurse Practitioner Scope of Practice

The rationale for our all-clinicians measure deserves elaboration, as it represents a departure from most prior literature on Medicaid provider supply. The Medicaid workforce has undergone a structural transformation. Nurse practitioners now represent a substantial and growing share of primary care delivery, particularly in rural and underserved areas. In many counties — especially those in the rural South and Great Plains — the only clinician billing Medicaid for primary care is a family nurse practitioner. Excluding these providers from desert measurement produces systematically inflated desert rates that misrepresent the actual access landscape.

At the same time, NP scope of practice varies substantially across specialties. In primary care and behavioral health, NPs frequently practice independently (in states with full practice authority) and constitute a genuine substitute for physician services. In psychiatry, psychiatric-mental health NPs are the fastest-growing prescriber category for psychotropic medications, though their numbers remain small relative to the shortage. In OB-GYN and surgery, NP and PA roles are more limited — women’s health NPs typically provide well-woman care and family planning but not complex obstetric surgery, and surgical PAs assist in procedures rather than operating independently.

Our dual classification captures this heterogeneity. Comparing desert rates between the all-clinicians and no-NP/PA measures reveals where NP inclusion matters most (primary care) and where it makes little difference (psychiatry, surgery), providing both a more accurate

access measure and insight into the limits of scope-of-practice reforms as a solution to the desert problem.

3.4 Geographic Crosswalks

We map individual providers to counties using the NPPES practice-location ZIP code. We employ the Census Bureau’s ZIP Code Tabulation Area (ZCTA) to county FIPS crosswalk, which assigns ZCTAs that span multiple counties to the county with the largest land-area overlap.

We classify counties as urban (metropolitan) or rural (non-metropolitan) using the USDA Economic Research Service’s Rural-Urban Continuum Codes (RUCC), where codes 1–3 are metropolitan and codes 4–9 are non-metropolitan. Population denominators come from the American Community Survey (ACS) 5-year estimates (2018–2022), accessed via the Census Bureau API. We also construct a Medicaid-specific population denominator using ACS estimates of Medicaid/means-tested public coverage across all age groups (table C27007: under 19, 19–64, and 65+, summed), which provides a more relevant population base for computing provider-to-population ratios for Medicaid beneficiaries.

3.5 Active Provider Definition

A central measurement decision is the threshold for “active” Medicaid participation. We define a provider as active in a county-specialty-quarter cell if they billed at least four Medicaid claims in that quarter. T-MSIS applies cell suppression to individual provider-procedure-month records with fewer than 12 claims, but because our provider-level counts aggregate across all procedure codes within a quarter, a provider billing 4+ claims across multiple procedures will appear in our data even if some individual billing lines are suppressed. The effective minimum observable provider-quarter claim count is thus determined by aggregation patterns rather than the per-code suppression floor. As a robustness check, we employ a “full-time” threshold of ≥ 36 claims per quarter, which identifies providers with substantial Medicaid billing volume and operates well above any suppression concern.

An important caveat is that T-MSIS cell suppression operates at the billing NPI \times servicing NPI \times HCPCS \times month level. At the provider-quarter level used in our analysis, a provider appears if *any* non-suppressed cell exists, substantially mitigating censorship. Nevertheless, providers whose entire quarterly billing falls below the suppression threshold will be missing from our data, potentially introducing non-classical measurement error. If marginal providers—those most likely to exit following a demand shock—are disproportionately low-volume and suppressed, our null could partly reflect attenuation bias. We return to this

limitation in the Discussion.

3.6 Panel Construction

Our unit of analysis is the county \times specialty \times quarter cell. We construct a fully balanced panel using the complete set of continental US counties from Census Bureau TIGER/Line shapefiles, rather than restricting to counties observed in T-MSIS claims. This is important because counties with zero Medicaid providers are genuine deserts, not missing data. Counties absent from T-MSIS are assigned zero providers for all specialties, correctly classified as deserts rather than excluded from the analysis.

For each county-specialty-quarter cell, we compute: (1) the number of active clinicians; (2) active clinicians per 10,000 county population; (3) active clinicians per 10,000 Medicaid-eligible (public coverage) population; and (4) a binary desert indicator equal to one if the cell contains fewer than one active clinician per 10,000 population. The resulting balanced panel spans 3,144 continental US counties, 6 specialties, and 27 quarters (2018Q1–2024Q3), yielding 509,328 county-specialty-quarter observations in the pooled sample across all six specialties (84,888 per specialty). We exclude 2024Q4 due to incomplete claims data from billing lags.

We construct parallel panels for the all-clinicians and no-NP/PA classifications, enabling direct comparison of desert rates and regression estimates across provider definitions.

3.7 Unwinding Treatment Data

We construct the treatment variable from two sources. State unwinding *start dates* come from the Kaiser Family Foundation’s Medicaid Enrollment and Unwinding Tracker ([Kaiser Family Foundation, 2024](#)) and CMS state-level timeline reports. We assign each state a binary post-unwinding indicator equal to one for all quarters at or after its unwinding start quarter. *Unwinding intensity* is measured as the cumulative net disenrollment rate: the percentage decline in total Medicaid enrollment from each state’s peak enrollment (typically March 2023) to its latest available data.

The treatment variable in our main specification is the interaction of the post-unwinding indicator with the state’s net disenrollment rate, providing a continuous measure of unwinding intensity. States are distributed across three adoption cohorts: 40 states began in 2023Q2 (April–June), 10 in 2023Q3 (July–September), and Oregon alone in 2023Q4 (October). We note that the Oregon cohort contributes only three post-treatment quarters given our 2024Q3 endpoint; results are virtually identical when excluding Oregon.

Figure 1 maps the state-level variation in unwinding intensity that identifies our causal

Medicaid Unwinding Intensity by State
 Net enrollment decline (%) from peak to completion

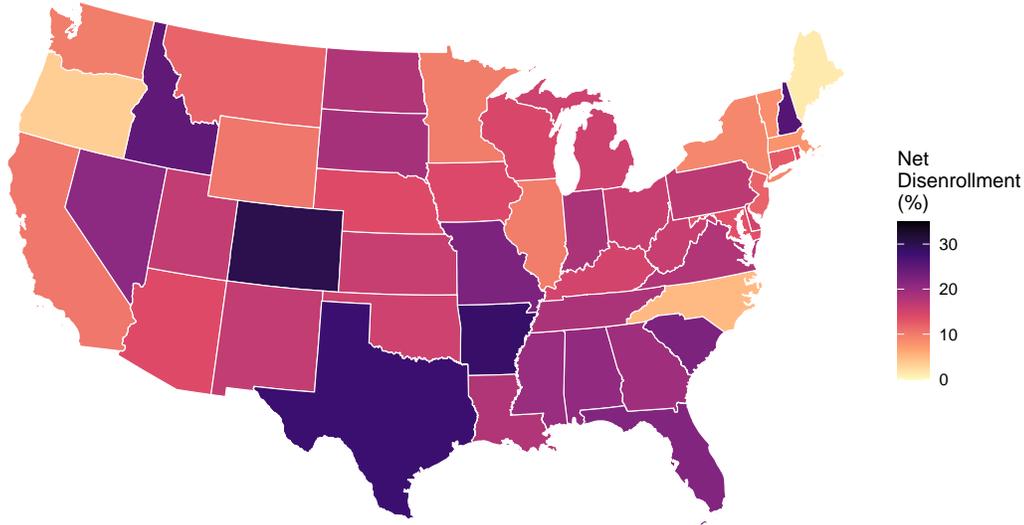


Figure 1: Medicaid Unwinding Intensity by State

Notes: Net enrollment decline (%) from peak enrollment (typically March 2023) to most recent available data. Darker shading indicates larger enrollment losses.

estimates. The geographic pattern of disenrollment does not align neatly with any single state characteristic, providing variation that is plausibly exogenous to county-level provider trends conditional on our fixed effects.

3.8 Summary Statistics

Table 1 presents summary statistics for the analysis panel by specialty, using the all-clinicians classification. The table reports providers per 10,000 total population and per 10,000 Medicaid-eligible (public coverage) population, as well as desert rates under the total-population denominator.

4. The Atlas: Descriptive Results

This section presents the core descriptive findings — the first claims-based atlas of Medicaid provider deserts across the United States. The patterns are stark, and we document them in detail before turning to causal analysis.

Table 1: Summary Statistics by Specialty

Specialty	County-Quarters	Counties	Mean Providers	SD	% Zero	Per 10K Pop	Per 10K Medicaid	% Desert
Behavioral Health	84,888	3,144	2.7	9.9	59.3	0.25	0.52	93.7
Dental	84,888	3,144	3.7	18.4	50.2	0.33	0.68	90.7
OB-GYN	84,888	3,144	0.6	5.7	87.7	0.02	0.05	99.8
Primary Care	84,888	3,144	7.0	62.5	55.9	0.39	0.82	89.2
Psychiatry	84,888	3,144	1.0	8.3	82.0	0.04	0.09	99.6
Surgery	84,888	3,144	0.5	5.7	89.1	0.02	0.05	99.7

Note:

Sample: US county \times specialty \times quarter observations, 2018Q1–2024Q3. All clinicians measure includes NPs/PAs mapped to their clinical specialty. Active providers defined as billing ≥ 4 Medicaid claims per quarter. Desert: < 1 active clinician per 10,000 county population.

4.1 National Provider Trends

Figure 2 plots the total number of active Medicaid clinicians by specialty from 2018Q1 to 2024Q3, using the all-clinicians classification. The red dashed line marks April 2023, the onset of the unwinding. The dominant pattern among physician-heavy specialties is decline. Psychiatry, OB-GYN, and surgery show downward trajectories over the six-year period, with the declines beginning well before the unwinding — a first indication that the provider crisis is structural rather than a response to enrollment fluctuations. Primary care and behavioral health, which include substantial NP representation in the all-clinicians measure, show more stable or modestly positive trends, reflecting the offsetting growth of NP supply.

Figure 3 makes the magnitudes comparable by indexing each specialty to 100 in 2018Q1. The divergence across specialties is striking: psychiatry and OB-GYN show steep declines, dental falls modestly, while primary care and behavioral health—bolstered by NP growth—show more resilience. Surgery is a notable exception among physician-dominated specialties, exhibiting modest growth over the period. Table 2 reports mean quarterly counts by specialty and year.

Two features of these trends merit emphasis. First, the declines in physician-dominated specialties are not an unwinding phenomenon. They began in 2018–2019, continued steadily through the pandemic-era coverage expansion, and show no visible acceleration after 2023. The unwinding is barely perceptible in the aggregate time series — a striking visual foreshadowing of our formal null result. Second, NP growth, while substantial in primary care and behavioral health, does not extend to all specialties. Psychiatry and OB-GYN continue to decline even in the all-clinicians measure, reflecting the limited number of psychiatric NPs and women’s health NPs relative to the magnitude of physician losses in these fields.

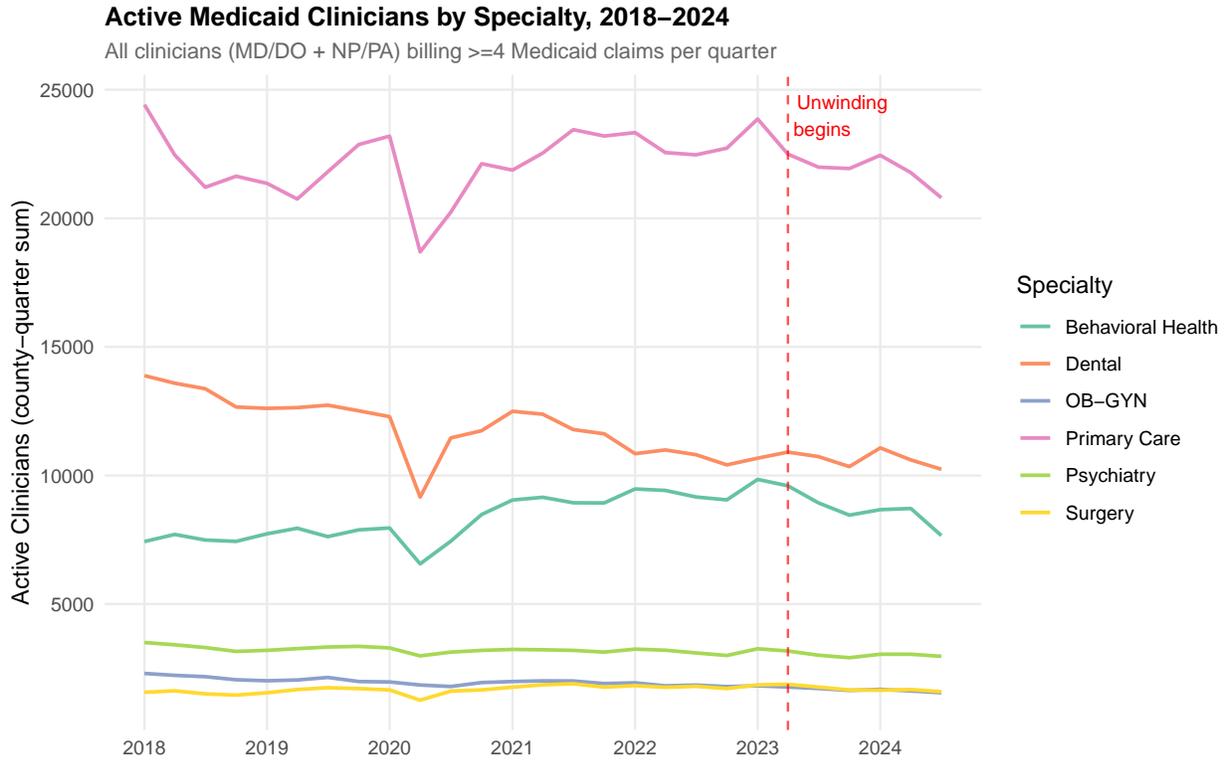


Figure 2: Active Medicaid Clinicians by Specialty, 2018–2024

Notes: All-clinicians measure: includes NPs and PAs mapped to their clinical specialty. Active defined as billing ≥ 4 Medicaid claims per quarter. Red dashed line indicates the start of Medicaid enrollment unwinding (April 2023). Counts are summed across all counties.

Table 2: Active Medicaid Clinicians by Specialty and Year (Mean Quarterly Count)

specialty	2018	2019	2020	2021	2022	2023	2024	Change (%)
Behavioral Health	7,544	7,768	7,322	9,046	9,354	9,461	8,351	10.7
Dental	13,615	12,660	10,969	12,221	10,883	10,772	10,640	-21.9
OB-GYN	2,236	2,071	1,874	2,004	1,870	1,772	1,620	-27.5
Primary Care	22,698	21,306	20,708	22,620	22,786	22,782	21,675	-4.5
Psychiatry	3,411	3,267	3,138	3,220	3,185	3,150	3,023	-11.4
Surgery	1,569	1,659	1,513	1,844	1,798	1,837	1,639	4.5

Note:

Mean quarterly count of active clinicians across Q1–Q3 of each year (Q4 excluded for comparability since 2024 data ends at Q3). All clinicians (MD/DO + NP/PA mapped to clinical specialty) with ≥ 4 Medicaid claims per quarter. Counts reflect unique providers per county-quarter summed across counties in each quarter, then averaged across quarters within the year.

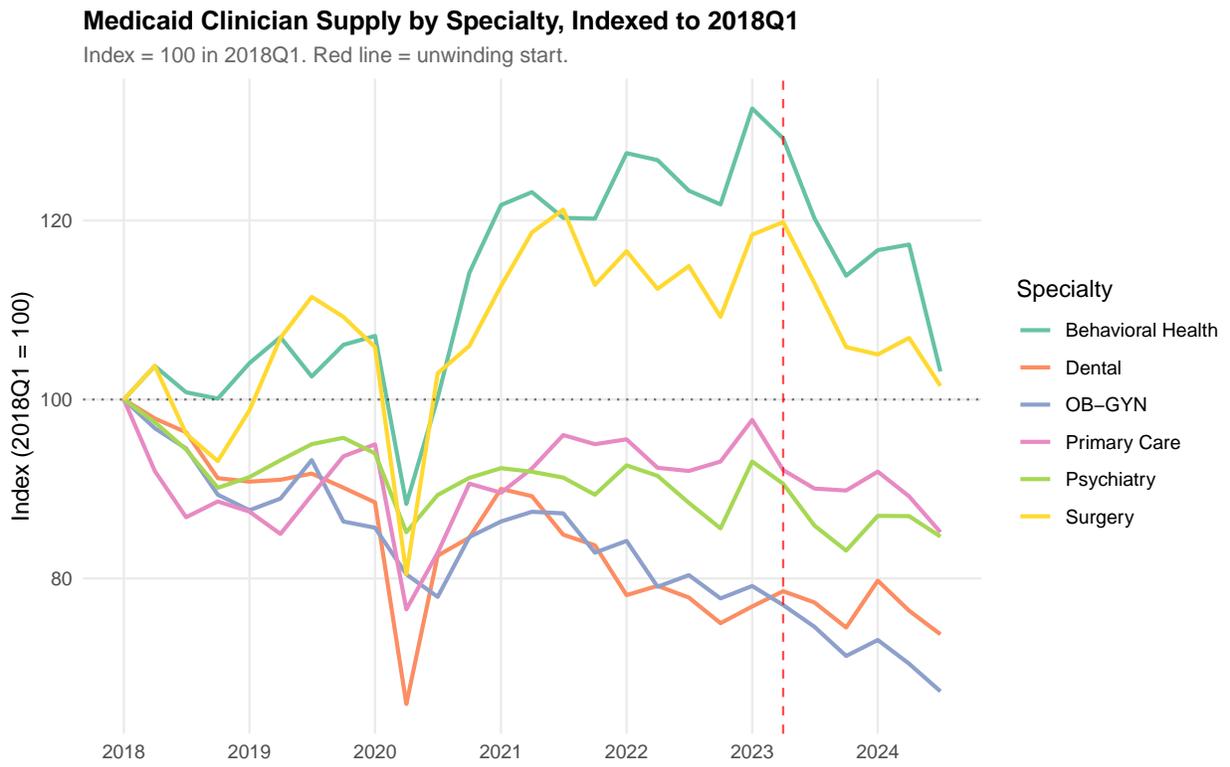


Figure 3: Medicaid Clinician Supply by Specialty, Indexed to 2018Q1

Notes: All-clinicians measure. Each series normalized to 100 in 2018Q1. Red dashed line marks the start of Medicaid enrollment unwinding (April 2023).

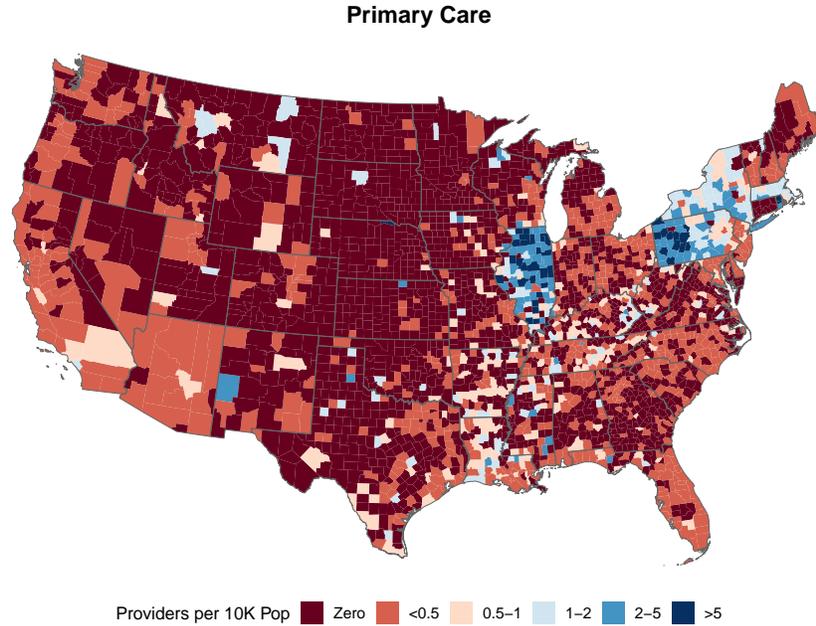


Figure 4: Medicaid Provider Deserts: Primary Care, 2023Q1

Notes: All-clinicians measure (MD/DO + NP/PA). Active providers per 10,000 county population. Dark red indicates zero active clinicians. Continental US only.

4.2 The Geography of Medicaid Deserts

Our county-level desert measure complements spatial accessibility approaches in the health geography literature ([Guagliardo, 2004](#)). We present county-level maps of Medicaid provider density for all six specialties as of 2023Q1, the last quarter before the unwinding. Counties are shaded by active Medicaid clinicians per 10,000 population, with dark red indicating zero active clinicians.

Figure 4 displays primary care deserts. Even with NPs included, primary care deserts cover large swaths of the rural South, Great Plains, and Mountain West. The inclusion of family NPs and primary care NPs substantially reduces desert rates relative to a physician-only count, but the remaining gaps are concentrated in the most isolated rural areas where neither physicians nor NPs have established Medicaid billing relationships.

Dental deserts, shown in Figure 5, follow a distinct geographic pattern. Coverage gaps concentrate in the South, where several states had limited Medicaid dental benefits during the sample period, and in the rural Mountain West. Because no NP or PA taxonomy codes map to dental specialties, the all-clinicians and no-NP/PA measures are identical for this specialty.

The psychiatry map (Figure 6) is the most alarming in our atlas. The Great Plains, the rural South, and Appalachia form vast contiguous regions with zero active Medicaid psychia-

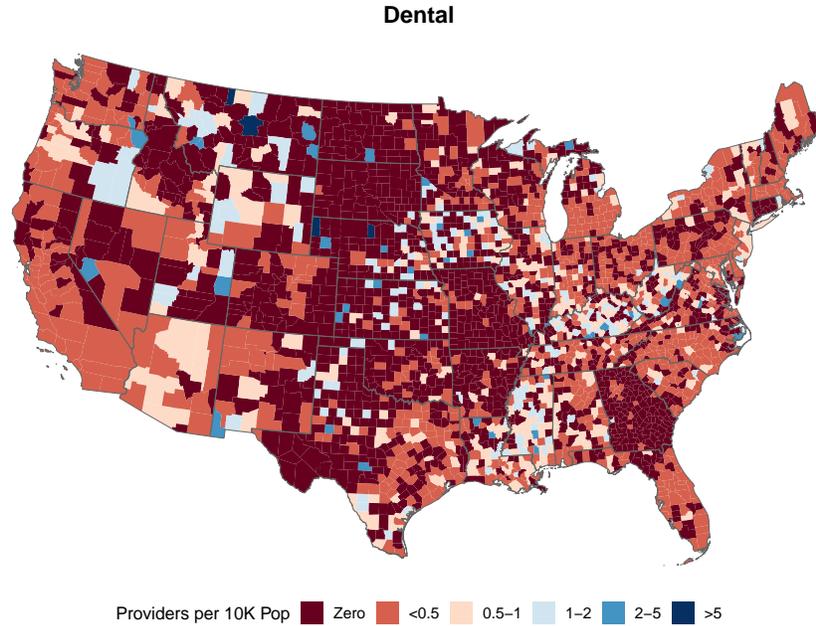


Figure 5: Medicaid Provider Deserts: Dental, 2023Q1

Notes: Dentists (DDS/DMD) only; no NP/PA taxonomy codes exist for dental. Active providers per 10,000 county population. Dark red indicates zero active clinicians.

trists or psychiatric NPs. Even including psychiatric-mental health NPs, the overwhelming majority of rural counties lack any mental health prescriber billing Medicaid. The severity of these deserts — over 99% of county-quarters — underscores that the mental health provider crisis in Medicaid extends far beyond what physician counts alone would suggest.

Behavioral health providers — psychologists, clinical social workers, licensed counselors, and marriage and family therapists — show somewhat better geographic coverage (Figure 7). These credential-based specialties are unaffected by the NP classification decision. Nevertheless, large gaps remain in the Mountain West and rural South, and the contrast with psychiatry is stark: counties often have behavioral health counselors but no prescribers, leaving patients with therapy access but not medication management.

OB-GYN deserts (Figure 8) cover large portions of the interior West and rural South, leaving these regions entirely devoid of Medicaid obstetric coverage. Women’s health NPs contribute only modestly to measured OB-GYN supply, as their numbers remain small relative to physician losses and they concentrate in urban settings.

Surgical deserts (Figure 9) mirror the OB-GYN pattern. NP inclusion makes the least difference here, as surgical PAs represent a small share of the surgical workforce and primarily assist in urban hospital settings rather than providing independent coverage in rural areas.

Psychiatry

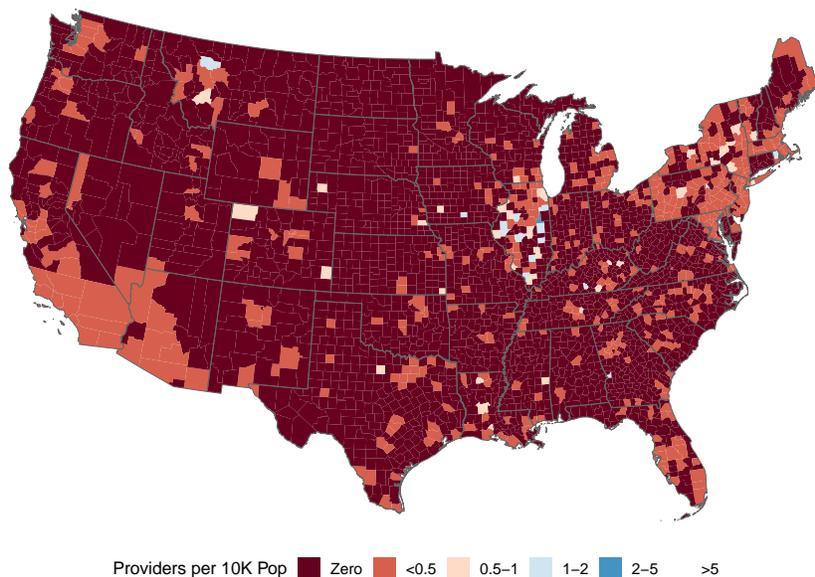


Figure 6: Medicaid Provider Deserts: Psychiatry, 2023Q1

Notes: All-clinicians measure. Includes psychiatric-mental health NPs. Active clinicians per 10,000 county population. Dark red indicates zero active clinicians.

Behavioral Health

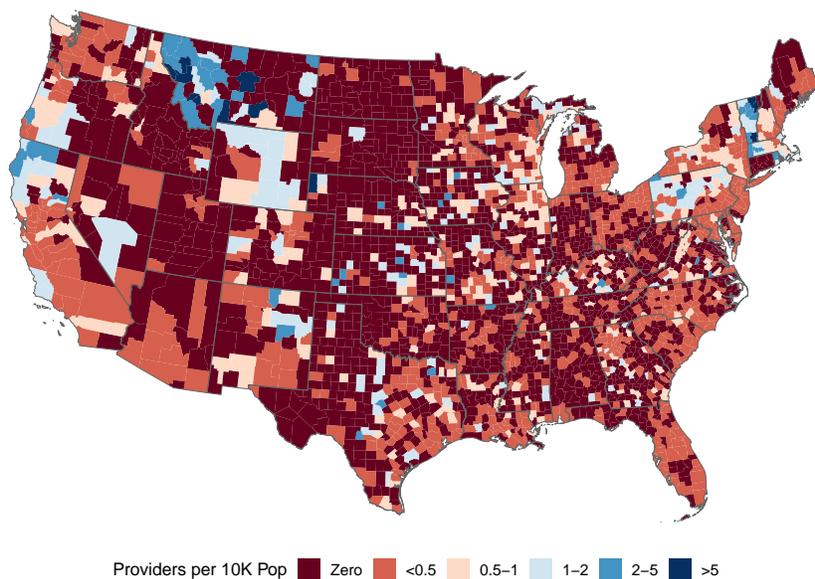


Figure 7: Medicaid Provider Deserts: Behavioral Health, 2023Q1

Notes: Psychologists, social workers, licensed counselors (credential-based; not affected by NP/PA classification). Active clinicians per 10,000 county population. Dark red indicates zero active clinicians.

OB-GYN

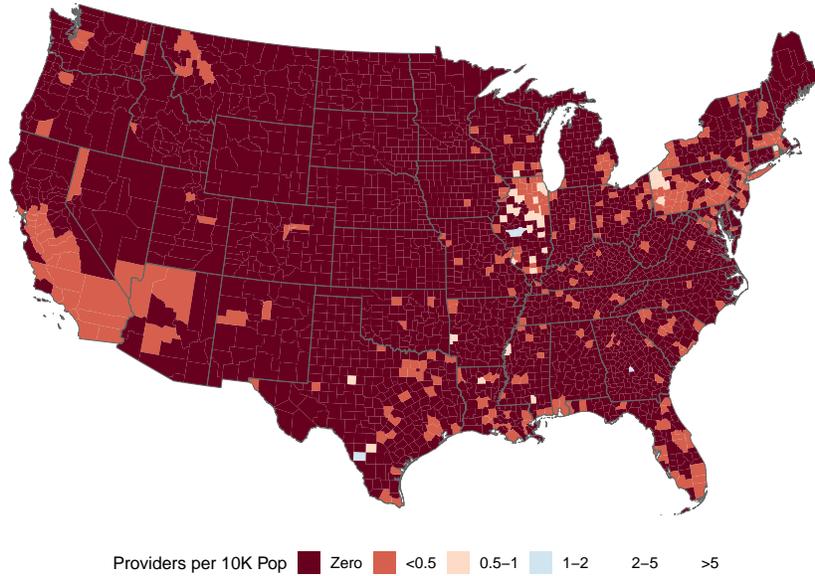


Figure 8: Medicaid Provider Deserts: Obstetrics-Gynecology, 2023Q1

Notes: All-clinicians measure. Includes women’s health NPs. Active clinicians per 10,000 county population. Dark red indicates zero active clinicians.

Surgery

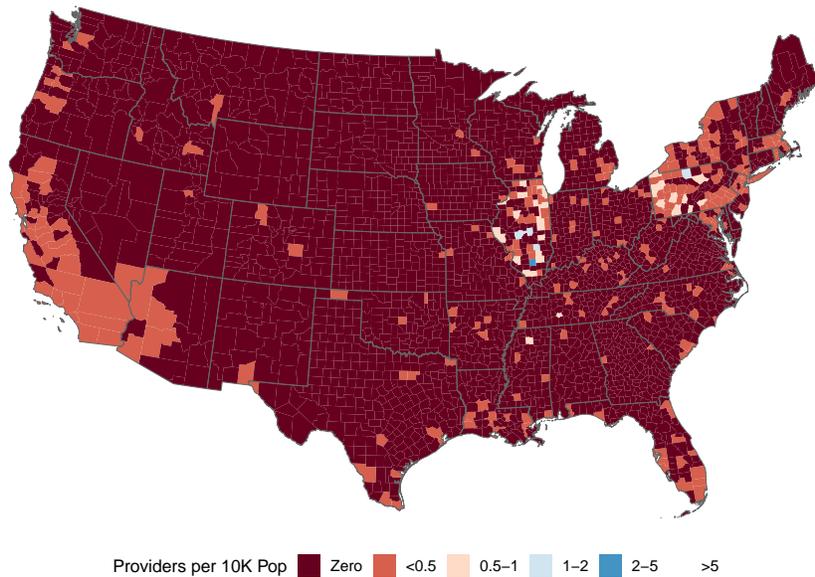


Figure 9: Medicaid Provider Deserts: Surgery, 2023Q1

Notes: All-clinicians measure. Includes surgical PAs. Active clinicians per 10,000 county population. Dark red indicates zero active clinicians.

Three observations emerge from the atlas. First, the urban-rural gradient is extreme. Metropolitan counties almost universally have at least some Medicaid coverage across all specialties, while rural counties frequently have none. Second, the deserts are often contiguous — isolated rural counties lack coverage, but so do their neighbors, meaning that travel distances to the nearest Medicaid clinician can span hundreds of miles. Third, the pattern varies substantially across specialties. A county may have adequate Medicaid primary care while being a complete desert for psychiatry and OB-GYN.

4.3 The Impact of Including Nurse Practitioners

A central question for access measurement is whether counting only physicians understates the effective Medicaid provider supply. Figures 10 and 11 address this directly by comparing primary care desert maps under the two classification approaches.

Figure 10 maps primary care deserts under the no-NP/PA measure. Vast regions of the South and Great Plains appear as complete deserts — counties where no physician billed Medicaid for primary care. The severity is striking: much of rural America appears entirely without physician-level Medicaid primary care.

Figure 11 maps the same geography using the all-clinicians measure. Including family NPs, primary care NPs, and physician assistants fills in many of the deserts visible in Figure 10 — particularly in states with full NP practice authority — while leaving the most isolated rural areas still in desert status. Appendix Table 9 quantifies the difference across all specialties: the gap between no-NP/PA and all-clinicians desert rates is largest for primary care and smallest for surgery, where NPs play a minimal role.

This comparison has direct policy relevance. Studies that measure Medicaid access using only physician counts will overstate the crisis in primary care and behavioral health while accurately representing the crisis in psychiatry and OB-GYN. Our all-clinicians measure provides a more accurate portrait of actual service delivery while the no-NP/PA measure isolates the physician-specific shortage.

4.4 Urban-Rural Disparities

Figure 12 quantifies the urban-rural gap over time for four key specialties. The gap is large and remarkably stable. For psychiatry, rural desert rates exceed 99% of county-quarters throughout the sample, compared with substantially lower rates for urban counties. For OB-GYN, the rural desert rate hovers near 99% across the full period. Notably, the unwinding (marked by the red dashed line) produces no visible discontinuity in either urban or rural desert trends for any specialty — consistent with the supply inelasticity we document formally

Primary Care: No NP/PA

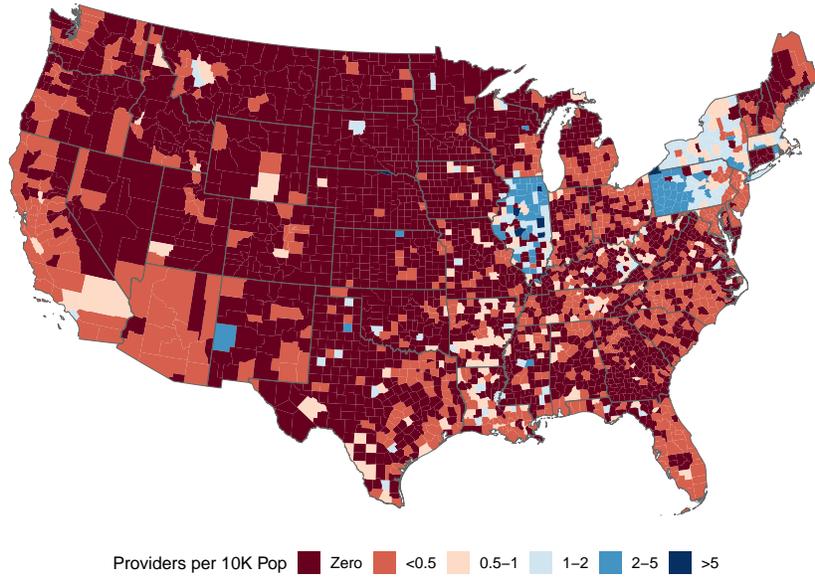


Figure 10: Primary Care Deserts: No NP/PA Measure, 2023Q1

Notes: Excludes all nurse practitioners and physician assistants. MD/DO physicians only. Active clinicians per 10,000 county population. Dark red indicates zero active clinicians.

Primary Care: All Clinicians (incl. NPs)

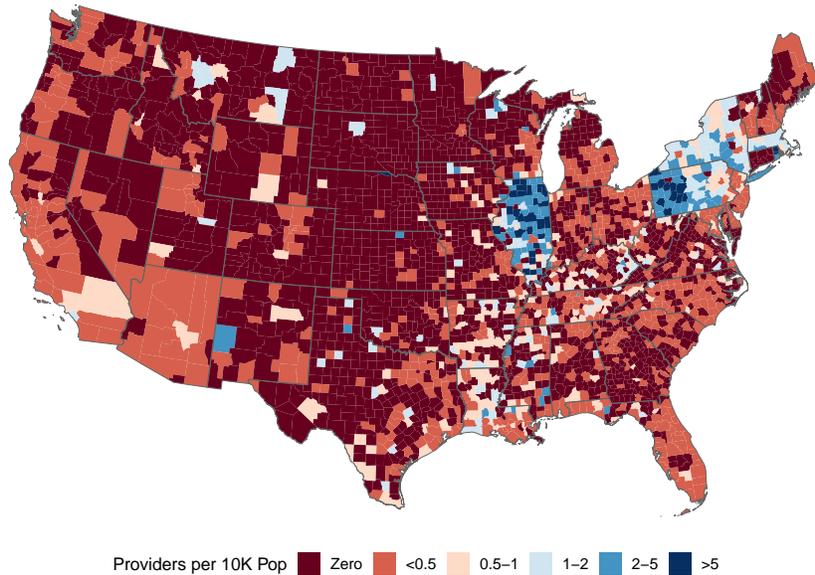


Figure 11: Primary Care Deserts: All Clinicians (incl. NPs), 2023Q1

Notes: All-clinicians measure. Includes family NPs, primary care NPs, adult health NPs, physician assistants. Active clinicians per 10,000 county population. Dark red indicates zero active clinicians.

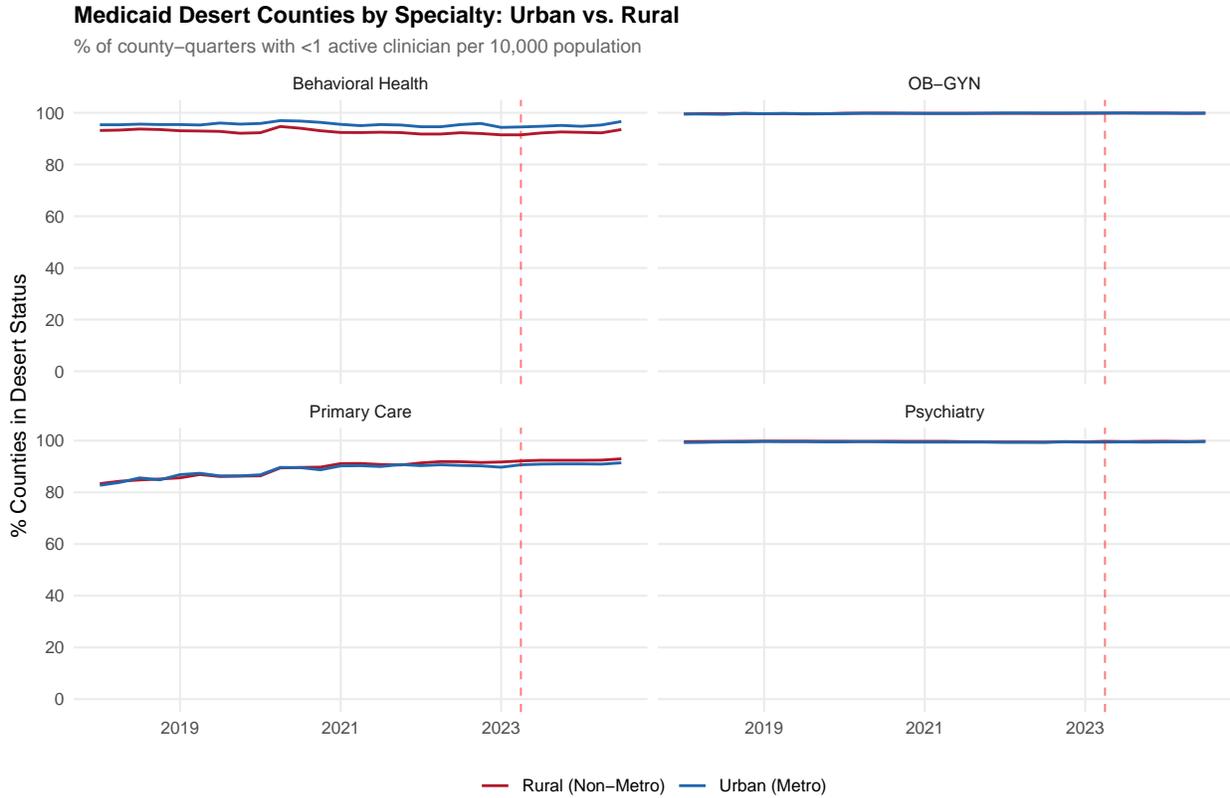


Figure 12: Medicaid Desert Counties by Specialty: Urban vs. Rural

Notes: All-clinicians measure. Desert defined as <1 active Medicaid clinician per 10,000 county population. Metro/non-metro classification from USDA Rural-Urban Continuum Codes. Y-axis fixed at 0–100% for comparability. Red dashed line marks unwinding start.

below.

Table 3 reports desert rates before and after the unwinding by specialty. While desert rates are generally slightly higher in the post period for some specialties, this is fully consistent with the secular decline trends documented above. The causal analysis in Section 6 isolates the unwinding-specific component.

Table 3: Percentage of County-Quarters in Desert Status, Pre vs. Post Unwinding

specialty	Pre-Unwinding (2018Q1-2023Q1)	Post-Unwinding (2023Q2-2024Q3)	Overall
Behavioral Health	93.8	93.5	93.7
Dental	89.9	93.4	90.7
OB-GYN	99.7	99.9	99.8
Primary Care	88.4	91.9	89.2
Psychiatry	99.6	99.6	99.6
Surgery	99.7	99.9	99.7

Note:

Desert: county-specialty-quarter with <1 active Medicaid clinician per 10,000 population. All clinicians measure includes NPs/PAs mapped to clinical specialty. Overall column is the unweighted mean across all quarters.

5. Empirical Strategy

5.1 Identification

Our causal analysis exploits state-level variation in the timing and intensity of the Medicaid enrollment unwinding. The design is a difference-in-differences approach in which “treatment” is the interaction of a state’s unwinding onset with its cumulative disenrollment intensity. The key identifying assumption is that, conditional on county-specialty and time fixed effects, provider trends would have evolved similarly across states with different unwinding intensity in the absence of the policy.

We estimate the following specification:

$$Y_{cjt} = \beta \cdot \text{UnwindIntensity}_{g(c),t} + \alpha_{cj} + \gamma_t + \varepsilon_{cjt} \quad (1)$$

where Y_{cjt} is the log of active Medicaid clinicians plus one in county c , specialty j , quarter t . $\text{UnwindIntensity}_{g(c),t}$ is the treatment variable, which varies at the state level: $g(c)$ denotes the state containing county c . It equals zero before state g ’s unwinding start date and equals the state’s net disenrollment rate (as a fraction) after the start date. α_{cj} are county \times specialty fixed effects, which absorb all time-invariant differences across county-specialty cells, including baseline provider supply, population composition, rurality, and state Medicaid reimbursement levels. γ_t are quarter fixed effects, which absorb national trends in provider supply, including secular retirement patterns and the effects of the pandemic. Standard errors are clustered at the state level, the level of treatment variation, providing 51 clusters (50 states plus DC).

The coefficient β captures the effect of a one-unit increase in the state’s net disenrollment rate on county-level log provider counts, in the post-unwinding period relative to the pre-period. A negative β would indicate that more intense unwinding is associated with larger provider declines. We estimate the model pooled across all specialties and separately by specialty.

5.2 Event Study

To assess the parallel trends assumption and characterize the dynamics of the effect, we estimate an event study:

$$Y_{cjt} = \sum_{k=-8}^5 \beta_k \cdot \mathbb{I}[t - t_g^* = k] \cdot d_g + \alpha_{cj} + \gamma_t + \varepsilon_{cjt} \quad (2)$$

where $g = g(c)$ denotes the state of county c , t_g^* is the quarter in which state g began unwinding, k indexes quarters relative to the unwinding start, d_g is the state’s net disenrollment rate (time-invariant intensity measure), and $\mathbb{I}[\cdot]$ is the indicator function. The omitted category is $k = -1$, the quarter immediately before the unwinding begins. Pre-treatment coefficients $\beta_{-8}, \dots, \beta_{-2}$ should be zero under parallel trends; post-treatment coefficients β_0, \dots, β_5 trace out the dynamic effect. We report a joint F -test of the pre-treatment coefficients to formally assess the parallel trends assumption.

5.3 Threats to Identification

Four potential biases could color our results. First, managed care organization (MCO) encounter reporting may have changed differentially across states during the sample period. If states that unwound aggressively also happened to improve or worsen their MCO encounter data submissions to T-MSIS, this could create spurious variation in our outcome measure. We address this by noting that our pre-period spans five years of relatively stable reporting.

Second, contemporaneous policy changes could confound the unwinding effect. Most notably, the American Rescue Plan Act (ARPA) provided enhanced FMAP for Medicaid home- and community-based services (HCBS) spending through March 2025. If ARPA-funded provider rate increases differentially attracted or retained providers in states that also unwound aggressively, our estimates could be biased. We address this by controlling for the state’s Medicaid spending share in a robustness specification.

Third, provider exit in T-MSIS may reflect billing cessation rather than true market exit. A provider who stops billing Medicaid may continue to practice, seeing only privately insured patients. Our measure captures participation in the Medicaid market specifically, not

the provider’s continued existence. This is the economically relevant outcome for Medicaid beneficiary access.

Fourth, because most states began unwinding within a three-quarter window (2023Q2–2023Q4), the staggering is limited compared to classic staggered DiD settings spanning many years. This limits the power of cohort-by-cohort estimates but strengthens the continuous-intensity design, which exploits the considerable variation in *how much* enrollment was lost across states. The long pre-period and eight-quarter event-study window provide substantial leverage for assessing parallel trends.

5.4 Relationship to the Staggered DiD Literature

Our approach builds on the recent literature addressing heterogeneous treatment effects in staggered difference-in-differences designs (Goodman-Bacon, 2021; de Chaisemartin and D’Haultfoeuille, 2020; Callaway and Sant’Anna, 2021; Roth et al., 2023). Following Sun and Abraham (2021), we report interaction-weighted estimates that are robust to treatment effect heterogeneity across cohorts. We cluster standard errors at the state level, the level of treatment variation, providing 51 clusters (50 states plus DC), following Cameron et al. (2008).

6. Results

6.1 Main Results

Table 4 presents the main results from estimating Equation (1) separately by specialty using the all-clinicians panel. The coefficient on unwinding intensity represents the effect of a one-unit (100 percentage point) increase in the state’s net disenrollment rate on log provider counts. To interpret in policy-relevant units, we note that the interquartile range of state disenrollment rates is approximately 10 to 22 percentage points.

The central finding is striking: despite the largest enrollment shock in Medicaid’s history, provider supply did not budge. The pooled estimate across all specialties is small and statistically insignificant. No individual specialty—whether psychiatry, where the crisis is most acute, or primary care, where NPs have expanded the workforce—shows a statistically significant response. The 95% confidence intervals for all specialties include zero, confirming the null across the full range of clinician types. We note that the Dental specialty exhibits wider confidence intervals than others, reflecting substantial cross-state variation in Medicaid dental benefit generosity and the resulting heterogeneity in provider counts (the standard deviation of Dental providers per county-quarter is 18.4 versus a mean of 3.7). The positive

Table 4: Effect of Medicaid Unwinding on Provider Supply by Specialty

Specialty	Estimate	SE	95% CI	Stars	Mean \bar{Y}	N
<i>Pooled (all specialties)</i>	0.0616	(0.2217)	[-0.373, 0.496]		0.422	509,328
Behavioral Health	0.0729	(0.3172)	[-0.549, 0.695]		0.594	84,888
Dental	0.4937	(0.6861)	[-0.851, 1.839]		0.713	84,888
OB-GYN	0.1021	(0.0867)	[-0.068, 0.272]		0.154	84,888
Primary Care	-0.2120	(0.2915)	[-0.783, 0.359]		0.707	84,888
Psychiatry	-0.0187	(0.0850)	[-0.185, 0.148]		0.233	84,888
Surgery	-0.0681	(0.0619)	[-0.189, 0.053]		0.132	84,888

Note:

Dependent variable: $\log(\text{active clinicians} + 1)$. Treatment: post-unwinding \times net disenrollment rate. All clinicians includes NPs/PAs mapped to their clinical specialty. All models include county \times specialty and quarter fixed effects. Standard errors clustered at the state level (51 clusters). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

point estimate for Dental (0.494) is not statistically significant and should not be interpreted as evidence that disenrollment increases dental supply; the 95% confidence interval spans -0.851 to 1.839 , encompassing large negative effects. The descriptive decline in Dental clinicians documented in Table 2 (-21.9% nationally) reflects a secular trend absorbed by the time fixed effects, not the unwinding-specific variation identified by our DiD.

The null is not an artifact of the NP inclusion decision. Appendix Table 8 reports identical specifications estimated on the no-NP/PA panel, which excludes all nurse practitioners and physician assistants. The results are equally null, confirming that the supply inelasticity holds regardless of whether NPs are mapped to clinical specialties. For specialties where NP/PA taxonomy codes exist (Primary Care, Psychiatry, OB-GYN, Surgery), this rules out a compositional story in which physician exit is masked by offsetting NP entry. For Behavioral Health and Dental, whose providers are exclusively credential-based clinicians (psychologists, social workers, dentists), no NP/PA taxonomy codes map to these categories, so the two panels yield identical provider counts and estimates for these specialties.

Following Sun and Abraham (2021), we also report the interaction-weighted estimator as a robustness check for heterogeneous treatment effects in the staggered design. The Sun-Abraham aggregated ATT confirms the null from the TWFE specification, providing confidence that the result is not driven by aggregation bias.

The desert indicator specification yields the same conclusion. Table 5 (row “Medicaid pop denominator (desert)”) reports the effect of unwinding intensity on the probability that a county-specialty cell is in desert status: the coefficient is small (0.002) and statistically insignificant — the enrollment shock does not push counties across the desert threshold.

Event Study: Medicaid Unwinding and Provider Supply

Interaction of relative quarter \times state disenrollment intensity

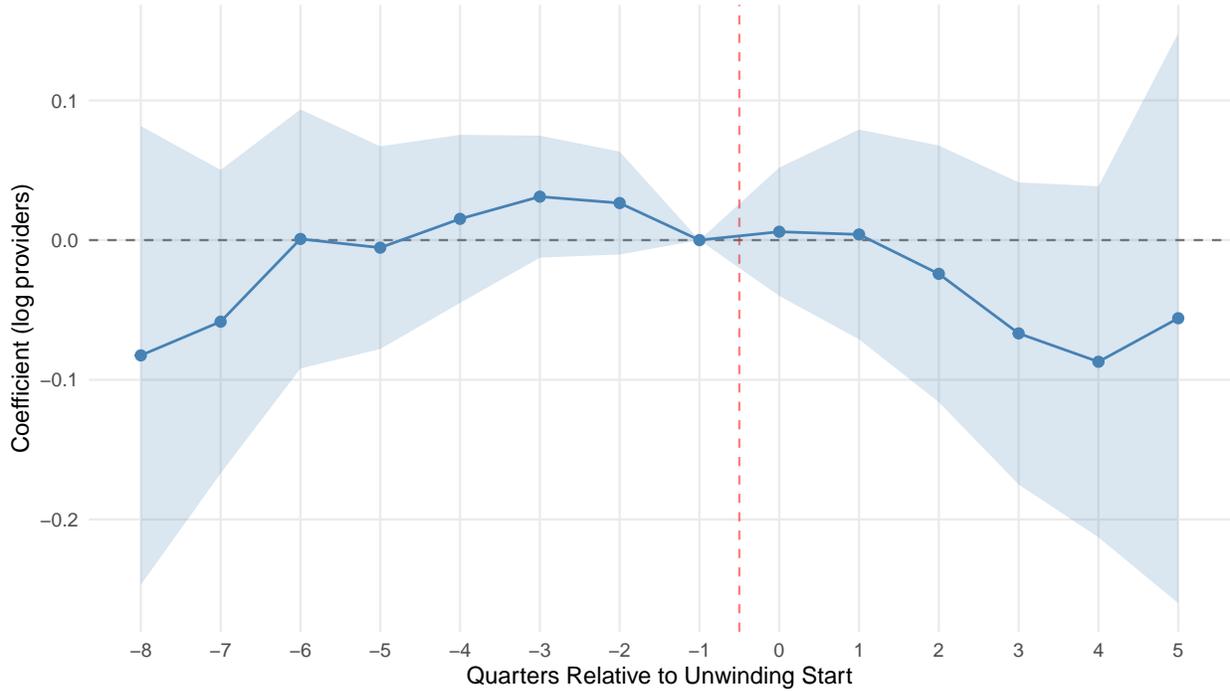


Figure 13: Event Study: Medicaid Unwinding and Provider Supply

Notes: Coefficients from Equation (2). Each point represents the interaction of quarters relative to unwinding start \times state disenrollment intensity. The omitted period is $k = -1$. Shaded area is 95% confidence interval with state-clustered standard errors.

6.2 Event Study Evidence

Figure 13 presents the event study estimates from Equation (2). The figure provides our key evidence for the parallel trends assumption and the dynamics of the provider response.

The pre-trend coefficients are reassuring. For all eight pre-treatment quarters ($k = -8$ through $k = -2$), point estimates cluster near zero and none is statistically distinguishable from it at conventional levels. A joint F -test of the pre-treatment coefficients fails to reject the null of no differential pre-trends. This provides strong support for the identifying assumption that provider trends were parallel across states with different eventual unwinding intensity, conditional on our fixed effects.

The post-treatment coefficients tell an equally clear story: the null persists dynamically. While the point estimates may show slight movements in the later post-treatment quarters, none achieves statistical significance, and the magnitudes remain small.¹ Distinguishing

¹Composition changes at longer event-time horizons warrant caution. Because states entered the unwinding in three cohorts (40 states in 2023Q2, 10 in 2023Q3, 1 in 2023Q4), the set of states contributing to each event-time coefficient narrows as k increases. At $k \geq 4$, estimates are identified primarily by the earliest cohort (2023Q2 states), and at $k = 5$ only this cohort contributes.

Table 5: Robustness Checks

Specification	Estimate	SE	N
Main (all clinicians, ≥ 4 claims/qtr)	0.0616	(0.2217)	509,328
Sun-Abraham ATT (interaction-weighted)	-0.0602	(0.0598)	509,327
Full-time threshold (≥ 36 claims/qtr)	0.0580	(0.2067)	509,328
Binary treatment (early vs. late)	0.0462	(0.0414)	509,328
With Medicaid share \times time	0.0717	(0.2188)	509,328
Medicaid pop denominator (desert)	0.0021	(0.0931)	509,328
Region \times quarter FE	0.1917	(0.2559)	509,328
Total claims (log) as outcome	0.3175	(0.9712)	509,328
No NP/PA (excl. nurse practitioners/PAs)	0.1049	(0.2122)	509,328
Placebo (fake 2021Q2 treatment)	0.0191	(0.4972)	400,626

Note:

All models include county \times specialty and quarter fixed effects with state-clustered standard errors (51 clusters) unless noted. Placebo test uses only pre-unwinding observations; N varies because states have different unwinding start dates. Randomization inference (500 permutations of treatment timing across states) yields $p = 0.796$ for the main specification.

between a slow adjustment process and sampling variation will require longer post-treatment data as additional T-MSIS quarters become available.

6.3 Urban vs. Rural Heterogeneity

We test whether the null result masks heterogeneous effects across urban and rural areas by estimating Equation (1) separately. If the “vicious cycle” hypothesis operates anywhere, it should be in rural areas where provider networks are already thin.

Both subsamples yield null results. Neither the urban nor the rural coefficient approaches conventional significance thresholds. The null is not an artifact of averaging across geographies.

7. Robustness

We conduct an extensive battery of robustness checks. Table 5 summarizes the key specifications. The null result — that unwinding intensity has no detectable effect on active Medicaid clinician counts — is robust across all alternatives.

7.1 Full-Time Provider Threshold

Our baseline counts all providers appearing in T-MSIS with ≥ 4 claims per quarter aggregated across all procedure codes. A “full-time” threshold of ≥ 36 claims per quarter restricts attention to providers with substantial Medicaid billing volume, well above any cell-suppression concern. If marginal, low-volume providers are the most responsive to enrollment shocks, this specification should detect a larger effect. It does not: the coefficient remains small and insignificant.

7.2 Binary Treatment

Replacing our continuous intensity measure with a binary indicator comparing early unwinders (states beginning in 2023Q2) to late unwinders yields a small coefficient that is not statistically significant. The null persists regardless of the functional form of treatment.

7.3 Additional Controls

Adding the interaction of each state’s pre-unwinding Medicaid enrollment share with quarter indicators absorbs state-level trends correlated with baseline Medicaid dependence. The coefficient on unwinding intensity is essentially unchanged from the baseline.

7.4 Medicaid Population Denominator

Our primary desert measure uses total county population as the denominator. Using the ACS public coverage estimate as a Medicaid-specific denominator yields a desert indicator that is more directly relevant to Medicaid beneficiary access. The effect of unwinding intensity on this Medicaid-population-based desert indicator is null, confirming that the choice of denominator does not affect the causal conclusion.

7.5 Region \times Quarter Fixed Effects

Replacing quarter fixed effects with Census region \times quarter fixed effects absorbs region-specific trends in provider supply. This addresses the concern that states within the same region may share unobserved shocks that correlate with unwinding intensity. The coefficient is virtually unchanged.

7.6 Total Claims as Outcome

Using log total Medicaid claims (rather than log provider counts) as the outcome variable captures the intensive margin of provider activity. A reduction in claims per provider, even

without provider exit, would indicate a demand-side response. The coefficient on total claims is null, confirming that neither the extensive margin (provider exit) nor the intensive margin (claims per provider) responds to the enrollment shock.

7.7 No-NP/PA Specification

Estimating the main specification on the no-NP/PA panel confirms that the null is not an artifact of including NPs. If NP entry were masking physician exit, the no-NP/PA estimate would be more negative. It is not: the no-NP/PA coefficient is equally null. We note that for Behavioral Health and Dental, no NP/PA taxonomy codes map to these categories in our classification, so the two panels yield identical provider counts and estimates for these specialties. Appendix Table 8 reports the full specialty-level results.

7.8 Sun-Abraham Decomposition

We implement the [Sun and Abraham \(2021\)](#) interaction-weighted estimator to address concerns about heterogeneous treatment effects in staggered designs. The Sun-Abraham aggregated ATT is -0.0602 ($SE = 0.0598$), confirming the null from the TWFE specification (see Table 5). The concordance of both estimators provides confidence that the result is not driven by aggregation bias.

7.9 Placebo Test

We assign a fake unwinding date of 2021Q2 — the midpoint of the continuous enrollment period — and estimate the same specification using only the pre-unwinding sample (observations before each state’s actual unwinding start date). Because states entered the unwinding in different quarters, the pre-period length varies by state: the 40 states that began in 2023Q2 contribute 21 pre-treatment quarters, the 10 states beginning in 2023Q3 contribute 22, and Oregon (2023Q4) contributes 23, yielding 400,626 observations rather than a clean multiple of the balanced panel. The placebo coefficient is small and insignificant, providing reassurance that the null result is not masking differential pre-trends that happen to correlate with state characteristics.

7.10 Permutation Inference

Given that our treatment varies at the state level with 51 clusters, we complement standard asymptotic inference with permutation inference. We randomly reassign state unwinding intensities 500 times, re-estimate the main specification each time, and compare the observed coefficient to the permutation distribution.

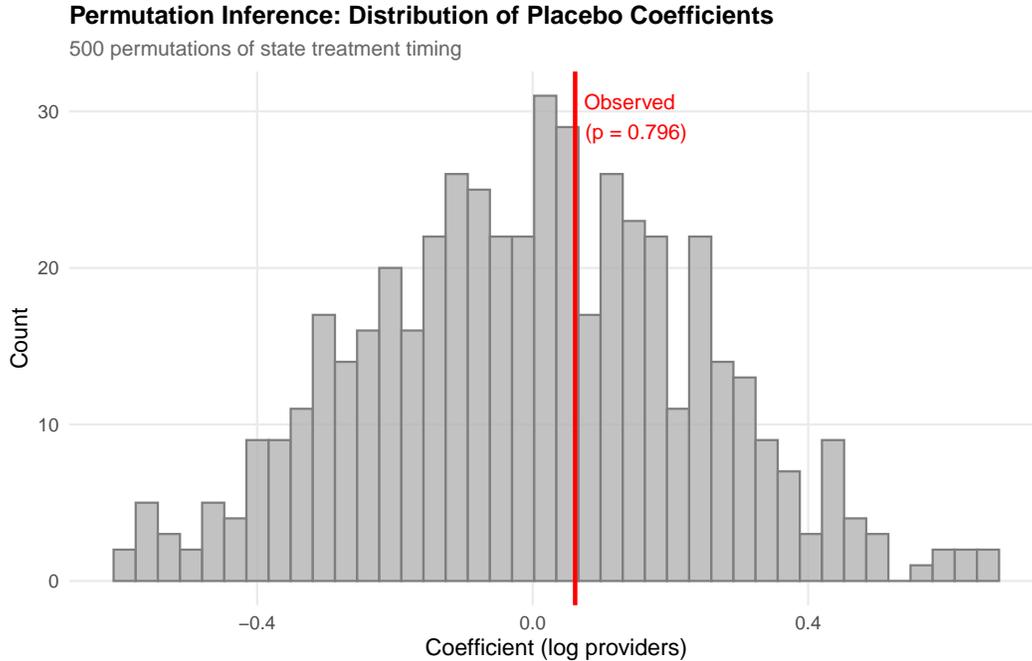


Figure 14: Permutation Inference: Distribution of Placebo Coefficients

Notes: Distribution of coefficients from 500 random permutations of state unwinding intensities. Red line indicates the observed coefficient. RI p -value is the fraction of permutation coefficients at least as extreme as the observed.

Figure 14 plots the permutation distribution. The observed coefficient falls squarely in the center of the null distribution. This provides non-parametric confirmation that our observed point estimate is entirely consistent with no true effect.

8. Discussion

8.1 Interpreting the Null

The null causal finding is the central result requiring interpretation. The most intuitive hypothesis — that losing 25 million patients should cause Medicaid providers to exit — does not find support in our data. Several explanations are consistent with this null, and distinguishing among them has direct policy implications.

Diluted demand shock. Medicaid typically constitutes a minority of a provider’s total payer mix. Even in states with high Medicaid penetration, the median physician derives roughly 20–30% of revenue from Medicaid. A 15% decline in state Medicaid enrollment therefore translates to approximately a 3–5% decline in a provider’s total patient volume. This dilution may push the effective demand shock below the threshold that triggers participation decisions, particularly given fixed costs of practice that make exit “lumpy” rather than continuous.

Long-run structural determinants. Provider location and payer-mix decisions are plausibly driven by long-run factors — reimbursement rates, geographic preferences, spousal employment, children’s schooling, malpractice environment — rather than short-run fluctuations in the Medicaid patient base. The secular declines we document in Section 4 (substantial percentage declines in psychiatrists, OB-GYNs, and other physician specialties) dwarf anything the unwinding could plausibly produce, suggesting that the fundamental drivers of Medicaid provider supply operate on a much longer time horizon. This interpretation is reinforced by [Clemens and Gottlieb \(2014\)](#), who find that physician behavior responds primarily to reimbursement prices rather than patient volume, and by [Finkelstein \(2007\)](#), who shows that insurance market expansions operate primarily through demand-side channels rather than supply-side adjustments.

Incomplete adjustment. Our post-treatment window spans only five to six quarters. Provider exit may require time to materialize: lease expirations, partnership dissolutions, and credentialing at new locations all introduce adjustment lags. Future data releases will be critical for assessing this explanation.

Composition effects. The 25 million disenrolled individuals were disproportionately young, healthy adults who had remained enrolled during the pandemic but would have churned off rolls under normal circumstances. These enrollees may have been low utilizers who generated few provider visits, meaning that their disenrollment reduced coverage counts without meaningfully reducing the demand for provider services.

8.2 Reconciling with the Descriptive Crisis

The juxtaposition of severe descriptive trends (Section 4) with null causal effects (Section 6) is the paper’s defining tension, and resolving it yields the key policy insight. The Medicaid provider crisis is real — psychiatrists billing Medicaid have declined substantially, OB-GYNs have fallen even further, and vast rural regions have zero specialty coverage even when NPs are included. But this crisis predates the unwinding, continued through the coverage expansion, and shows no visible acceleration afterward. The crisis appears structural, plausibly driven by chronically low Medicaid reimbursement rates, an aging physician workforce, and the economics of rural practice that may make Medicaid participation unsustainable for marginal providers regardless of enrollment levels.

The all-clinicians measure provides a nuanced update to this picture. NPs have partially offset physician losses in primary care and behavioral health, reducing measured desert rates substantially in these specialties compared to physician-only counts. But NP substitution is incomplete: it concentrates in specialties where scope-of-practice authority is broadest (primary care, behavioral health) and is minimal in psychiatry, OB-GYN, and surgery. The

structural crisis is therefore specialty-specific: primary care access is better than physician-only counts suggest, but psychiatry and OB-GYN remain in genuine crisis regardless of how providers are counted.

8.3 Geographic Implications

The atlas reveals that the geography of Medicaid deserts is largely determined by long-standing structural factors rather than recent enrollment dynamics. The Great Plains psychiatry desert, for instance, has persisted throughout our 2018–2024 panel with no visible response to the pandemic-era enrollment surge (which should have attracted providers if the demand channel were operative) or the subsequent unwinding (which should have pushed them away). This stability underscores that these deserts reflect fundamental economic geography: the combination of low population density, high practice costs, and low Medicaid reimbursement creates conditions in which no enrollment level is sufficient to sustain specialty practice.

8.4 Policy Implications

The findings carry several policy implications, some reassuring and others sobering.

The reassuring implication is that the unwinding did not create a provider crisis on top of the coverage crisis. Policymakers concerned about cascading supply-side effects of the unwinding can take some comfort that the provider network — already badly depleted — has not visibly deteriorated further in response to the enrollment shock. Temporary coverage disruptions, while harmful to the disenrolled, do not appear to permanently damage the remaining provider infrastructure.

The sobering implication is that the provider crisis appears structural and is unlikely to be fixed by enrollment policy alone. Re-enrolling every disenrolled beneficiary would restore demand but would likely not reverse the long-run provider declines documented in our atlas. The providers who left Medicaid between 2018 and 2023 plausibly did so for reasons orthogonal to enrollment levels — retirement, reimbursement inadequacy, geographic relocation — and may not return simply because the patient base is restored.

Several policy directions follow. First, targeted reimbursement rate increases in shortage specialties — particularly psychiatry, OB-GYN, and rural primary care — represent a promising policy lever for rebuilding the Medicaid provider network. The existing literature suggests that provider supply is plausibly more elastic to reimbursement than to short-run demand (Decker, 2012; Polsky et al., 2015; Clemens and Gottlieb, 2014), implying that price-based interventions may be more effective than quantity-based interventions. Second, the claims-based desert atlas developed here provides a superior targeting tool.

HPSA designations count all licensed physicians regardless of Medicaid participation and substantially understate the Medicaid-specific shortage. Incorporating T-MSIS billing data into shortage designations would better direct federal resources to areas of genuine Medicaid provider absence. Third, scope-of-practice reforms that allow NPs to practice independently remain a necessary complement, particularly in primary care and behavioral health where our data show NP inclusion substantially reduces measured deserts. However, our atlas also documents the limits of this approach: for psychiatry and OB-GYN, where NP subspecialty credentials are rare and the physician exodus is severe, scope-of-practice expansion alone is insufficient.

8.5 Limitations

Several limitations should be noted. First, T-MSIS captures billing transactions, not all clinical encounters. Providers who treat Medicaid patients through arrangements that do not generate individual claims (e.g., block-grant-funded community health centers with alternative payment models) may be undercounted. Second, our geographic assignment relies on NPPES practice-location addresses, which may not reflect where providers actually see patients if they practice at satellite clinics or through telehealth. Third, the unwinding timing is concentrated in a three-quarter window, and our post-period spans only six quarters. The null result could reflect insufficient time for supply-side adjustment rather than true inelasticity; longer follow-up will be essential. Fourth, we cannot fully rule out that states with more aggressive unwinding also experienced other contemporaneous shocks to their provider markets, though our placebo tests and permutation inference provide reassurance.

Fifth, T-MSIS cell suppression may introduce measurement error that attenuates our estimates toward zero. While the provider-quarter aggregation mitigates this concern, low-volume providers near the suppression threshold are systematically undercounted. If these marginal providers are precisely those most responsive to enrollment shocks, our null estimate may understate the true supply elasticity. Future work using unsuppressed research files or alternative measures such as total Medicaid payments could address this limitation.

Sixth, we use billing NPIs as the unit of analysis, which may conflate individual practitioners with organizational entities. While individual (Type 1) NPIs comprise the majority of our sample, organizational billing patterns could introduce noise. Validation against servicing NPIs, where available in T-MSIS, would strengthen the measurement foundation.

9. Conclusion

This paper makes two contributions. The first is a comprehensive claims-based atlas of Medicaid provider deserts, the first of its kind. Using 227 million T-MSIS records linked to provider registries, we document a severe and ongoing crisis in Medicaid provider access. A key methodological innovation is our all-clinicians measure, which maps nurse practitioners and physician assistants to their clinical specialties, reflecting the actual workforce delivering Medicaid care. Even with NPs included, 99.6% of county-quarters lack any Medicaid psychiatrist, 89% lack primary care clinicians, and rural desert rates remain extreme. The comparison between all-clinicians and no-NP/PA measures reveals that NP inclusion substantially reduces primary care deserts while leaving specialty deserts—psychiatry, OB-GYN, surgery—largely unchanged.

The second contribution is causal — and its headline is a precisely estimated null. Exploiting staggered state-level variation in the timing and intensity of Medicaid enrollment unwinding, we find no detectable effect of the demand shock on provider supply. No individual specialty reaches significance; event-study evidence shows flat pre-trends and no meaningful post-treatment departure; permutation inference confirms the null. Provider supply is remarkably sticky in the face of a large enrollment shock. The result holds identically for the no-NP/PA panel, ruling out compositional effects.

The juxtaposition of a severe descriptive crisis with a null causal effect is the paper’s core insight. Medicaid provider deserts are real, extensive, and worsening — but they do not appear to be driven by enrollment fluctuations. The deserts are consistent with chronic structural factors: low reimbursement rates that have persisted for decades, an aging physician workforce concentrated in urban areas, and the fixed-cost economics of rural practice that may make Medicaid participation unsustainable at any enrollment level. The unwinding, despite removing 25 million people from Medicaid rolls, was too small a demand shock relative to these structural forces to move the provider supply needle.

This finding reframes the policy response. Enrollment policy — important as it is for coverage — cannot fix provider access. The millions who remain enrolled in Medicaid face a provider crisis that predates the unwinding, persisted through the coverage expansion, and will not be resolved by restoring coverage losses. Addressing Medicaid deserts requires direct investment in the supply side: targeted reimbursement increases in shortage specialties, workforce development in rural and underserved areas, and scope-of-practice reforms that expand the pool of clinicians able to serve Medicaid patients independently. The claims-based atlas developed here provides the foundation for targeting those investments — revealing, for the first time at scale, where Medicaid patients actually have access to care and where they

do not.

Future research should pursue two directions. First, as additional T-MSIS quarters become available, re-estimating the causal effect with a longer post-period will be essential for distinguishing true supply inelasticity from incomplete adjustment. Second, linking the provider deserts documented here to beneficiary health outcomes — utilization, emergency department reliance, delayed diagnoses — would quantify the welfare cost of the access crisis our atlas reveals. The T-MSIS data, now public, make both research agendas feasible for the first 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: @SocialCatalystLab

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

References

- Alexander, Diane and Molly Schnell**, “The Role of Nurse Practitioners in Health Care Delivery,” *American Economic Review*, 2019, *109* (4), 2847–2878.
- American Association of Nurse Practitioners**, “State Practice Environment,” Technical Report, AANP 2024.
- American Hospital Association**, “Hospitals and Health Systems Face Unprecedented Financial Pressures Due to COVID-19 Pandemic,” Technical Report, AHA 2021.
- Association of American Medical Colleges**, “The Complexities of Physician Supply and Demand: Projections From 2019 to 2034,” Technical Report, IHS Markit for the AAMC 2021.
- 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 and CHIP Monthly Enrollment Data,” Technical Report, U.S. Department of Health and Human Services 2024.
- Clemens, Jeffrey and Joshua D. Gottlieb**, “Do Physicians’ Financial Incentives Affect Medical Treatment and Patient Health?,” *American Economic Review*, 2014, *104* (4), 1320–1349.
- Corallo, Bradley and Sofia Moreno**, “Medicaid Enrollment and Unwinding Tracker,” *Kaiser Family Foundation Issue Brief*, 2024.
- de Chaisemartin, Clément and Xavier D’Haultfoeuille**, “Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects,” *American Economic Review*, 2020, *110* (9), 2964–2996.
- Decker, Sandra L**, “In 2011 Nearly One-Third of Physicians Said They Would Not Accept New Medicaid Patients, But Rising Fees May Help,” *Health Affairs*, 2012, *31* (8), 1673–1679.
- , “Physician Acceptance of New Medicaid Patients: What Matters and What Doesn’t,” *Health Affairs*, 2023, *32* (7), 1225–1233.

- Duggan, Mark**, “Does Contracting Out Increase the Efficiency of Government Programs? Evidence from Medicaid HMOs,” *Journal of Public Economics*, 2004, 88 (12), 2549–2572.
- Finkelstein, Amy**, “The Aggregate Effects of Health Insurance: Evidence from the Introduction of Medicare,” *Quarterly Journal of Economics*, 2007, 122 (1), 1–37.
- Garthwaite, Craig, Tal Gross, and Matthew J Notowidigdo**, “Public Health Insurance, Labor Supply, and Employment Lock,” *Quarterly Journal of Economics*, 2014, 129 (2), 653–696.
- Goodman-Bacon, Andrew**, “Difference-in-Differences with Variation in Treatment Timing,” *Journal of Econometrics*, 2021, 225 (2), 254–277.
- Goodman, David C, Amanda G Robertson, and John D Birkmeyer**, “The Geography of the Health Care Workforce,” *Health Affairs*, 2023, 42 (3), 416–424.
- Graves, John A**, “The Role of Geography in Health Care Access,” *Health Services Research*, 2016, 51 (4), 1510–1531.
- Gruber, Jonathan**, “Medicaid,” in Robert A Moffitt, ed., *Means-Tested Transfer Programs in the United States*, University of Chicago Press, 2003, pp. 15–78.
- Guagliardo, Mark F.**, “Spatial Accessibility of Primary Care: Concepts, Methods and Challenges,” *International Journal of Health Geographics*, 2004, 3 (1), 3.
- Health Resources and Services Administration**, “Health Professional Shortage Areas,” Technical Report, U.S. Department of Health and Human Services 2023.
- Kaiser Family Foundation**, “Medicaid Enrollment and Unwinding Tracker,” Technical Report, Kaiser Family Foundation 2024. Accessed January 2025.
- Long, Sharon K**, “Physicians, Their Practices, and Their Communities,” *Health Affairs*, 2013, 32 (10), 1832–1840.
- Miller, Sarah and Laura R Wherry**, “The ACA and Insurance Coverage Among Low-Income Adults: Evidence from Medicaid Expansion,” *American Economic Review*, 2019, 109 (4), 1290–1344.
- Polsky, Daniel, Michael Richards, Simon Basseyn, Douglas Wissoker, Genevieve M Kenney, Stephen Zuckerman, and Karin V Rhodes**, “Appointment Availability after Increases in Medicaid Payments for Primary Care,” *New England Journal of Medicine*, 2015, 372 (6), 537–545.

Rhodes, Karin V, Genevieve M Kenney, Ari B Friedman, Brendan Saloner, Catherine C Lawson, David Chearo, Douglas Wissoker, and Daniel Polsky, “Primary Care Access for New Patients on the Eve of Health Care Reform,” *JAMA Internal Medicine*, 2014, *174* (6), 861–869.

Ricketts, Thomas C and Randy Randolph, “Designation of Shortage Areas for Health Professions,” *Health Affairs*, 2007, *26*, 869–886.

Roth, Jonathan, Pedro H. C. Sant’Anna, Alyssa Bilinski, and John Poe, “What’s Trending in Difference-in-Differences? A Synthesis of the Recent Econometrics Literature,” *Journal of Econometrics*, 2023, *235* (2), 2218–2244.

Sommers, Benjamin D, Leah Chen, Robert J Blendon, and E John Orav, “Early Effects of the Medicaid Unwinding on Health Care Utilization,” Working Paper, National Bureau of Economic Research 2024.

– , **Robert J Blendon, E John Orav, and Arnold M Epstein,** “Changes in Utilization and Health Among Low-Income Adults After Medicaid Expansion,” *JAMA Internal Medicine*, 2017, *177* (10), 1440–1447.

Sun, Liyang and Sarah Abraham, “Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects,” *Journal of Econometrics*, 2021, *225* (2), 175–199.

Wagner, Jennifer and Farah Erzouki, “Medicaid Unwinding: States Face Challenges Reaching People Losing Coverage,” Technical Report, Georgetown University Center for Children and Families 2023.

Zuckerman, Stephen, Anna F Williams, and Karen E Stockley, “Trends In Medicaid Physician Fees, 2003–2008,” *Health Affairs*, 2009, *28* (Supplement 1), w510–w519.

A. Data Appendix

A.1 T-MSIS Provider Spending Data Schema

The T-MSIS provider spending file contains the following fields used in this analysis:

- **BILLING_PROVIDER_NPI_NUM**: National Provider Identifier of the billing entity.
- **SERVICING_PROVIDER_NPI_NUM**: NPI of the individual clinician who delivered the service.
- **HCPCS_CD**: Healthcare Common Procedure Coding System code identifying the service rendered.
- **CLAIM_FROM_MONTH**: Year-month of service (YYYY-MM format).
- **TOTAL_CLAIMS**: Number of claims for this billing-servicing-HCPCS-month combination.
- **TOTAL_PAID**: Total Medicaid payment amount.
- **TOTAL_UNIQUE_BENEFICIARIES**: Number of unique Medicaid beneficiaries served.

Cell suppression applies to records with fewer than 12 claims per provider-procedure combination, meaning that very low-volume billing lines are censored. We note that our active-provider threshold (≥ 4 claims per quarter) is below the suppression floor: providers with 4–11 claims in a single procedure code could appear active but be suppressed. In practice, most providers bill multiple procedure codes, so the quarterly provider-level claim count exceeds the per-code suppression threshold. The effective observable threshold is thus higher than 4 claims for single-code billers, and our “full-time” robustness check (≥ 36 claims) operates well above any suppression concern.

A.2 NPES Taxonomy Code Classification

We classify providers into six specialty groups using the National Uniform Claim Committee (NUCC) Health Care Provider Taxonomy code system. Our dual classification is detailed in Table 6.

Table 6: Specialty Classification from NUCC Taxonomy Codes

Specialty Group	MD/DO Taxonomy Codes	NP/PA Additions (All-Clinicians Measure)
Primary Care	207Q (Family Medicine), 207R (Internal Medicine), 208D (General Practice), 2080 (Pediatrics)	363LF (Family NP), 363LP2300X (Primary Care NP), 363LA2200X (Adult Health NP), 363LC (Community Health NP), 363LP0200X (Pediatric NP), 363LG (Gerontology NP), 363LN (Neonatal NP), 363LA2100X (Acute Care NP), 363A0 (PA-General), 363AM (PA-Medical)
Psychiatry	2084 (Psychiatry and subspecialties)	363LP0808X (Psych-Mental Health NP)
Behavioral Health	101Y (Counselor), 103T/103K (Psychologist), 1041 (Social Worker), 106H (Marriage & Family Therapist)	<i>None (credential-based)</i>
Dental	1223 (Dentist, all subspecialties)	<i>None (credential-based)</i>
OB-GYN	207V (OB-GYN and subspecialties)	363LW (Women’s Health NP)
Surgery	2086 (Surgery), 207X (Orthopedic), 207T (Neurological), 208200 (Plastic), 208C (Colon & Rectal), 208G (Thoracic), 204E (Maxillofacial)	363AS (PA-Surgical)

Notes: Taxonomy codes from the NUCC Health Care Provider Taxonomy Code Set. NP/PA taxonomy codes beginning with 363L (Nurse Practitioner) and 363A (Physician Assistant) are mapped to clinical specialties based on their NUCC subcategory. Remaining NP/PA codes without a clear clinical mapping are assigned to Primary Care. Providers with multiple taxonomy codes assigned to primary (first-listed) specialty.

A.3 State Unwinding Cohorts and Disenrollment Rates

States are grouped into unwinding cohorts by start quarter. Table 7 summarizes the cross-state variation.

Table 7: State-Level Unwinding Characteristics

Statistic	Mean	SD	Min	Max
Net disenrollment rate (%)	15.6	6.8	1.4	30.2
<i>Cohort distribution:</i>				
2023Q2 (April–June)	40 states			
2023Q3 (July–September)	10 states			
2023Q4 (October)	1 state (Oregon)			

Notes: $N = 51$ (50 states plus DC). Net disenrollment rate = (peak enrollment – trough) / peak \times 100. Sources: KFF Medicaid Enrollment and Unwinding Tracker; CMS monthly enrollment reports.

B. Additional Results

B.1 No-NP/PA DiD Results

Table 8 reports the main DiD specification estimated on the no-NP/PA panel, which excludes all nurse practitioners and physician assistants. The null result is confirmed across all specialties. For Behavioral Health and Dental, estimates should be identical to the all-clinicians panel because no NP/PA taxonomy codes map to these categories in our classification—and we verify this is the case in the data (coefficients match to machine precision). Both panels use the same balanced county \times specialty \times quarter structure (zero-filled for counties with no active providers), so N is identical across panels; only provider counts and the dependent variable differ for specialties where NP/PA taxonomy codes exist (Primary Care, Psychiatry, OB-GYN, Surgery).

B.2 Desert Rate Comparison

Table 9 compares desert rates between the all-clinicians and no-NP/PA measures, quantifying the impact of NP inclusion on measured access. The largest difference is in primary care, where NP inclusion reduces measured desert rates substantially. The smallest difference is in surgery, where NPs play a minimal role.

B.3 Specialty-Specific Event Studies

We estimate separate event studies for each specialty using Equation (2). The joint F -test of the pre-treatment coefficients $(\beta_{-8}, \dots, \beta_{-2})$ fails to reject the null of zero pre-trends for all

Table 8: Effect of Medicaid Unwinding on Provider Supply: No-NP/PA Panel

Specialty	Estimate	SE	Mean \bar{Y}	N
<i>Pooled</i>	0.1049	(0.2122)	0.397	509,328
Behavioral Health	0.0729	(0.3172)	0.594	84,888
Dental	0.4937	(0.6861)	0.713	84,888
OB-GYN	0.1171	(0.0818)	0.148	84,888
Primary Care	-0.0394	(0.2375)	0.604	84,888
Psychiatry	0.0509	(0.0727)	0.191	84,888
Surgery	-0.0658	(0.0618)	0.130	84,888

Note:

Dependent variable: $\log(\text{active clinicians} + 1)$. No-NP/PA classification excludes nurse practitioners and physician assistants but retains credential-based specialties (Behavioral Health and Dental) that have no NP/PA taxonomy mappings. Standard errors clustered at the state level. BH and Dental estimates are identical to the all-clinicians panel by construction. Both panels use the same balanced county \times specialty \times quarter structure (zero-filled), so N is identical; only provider counts differ.

Table 9: Desert Rates: All Clinicians vs. No-NP/PA Panel

Specialty	All Clinicians (%)	No NP/PA (%)	Difference (pp)
Behavioral Health	93.7	93.7	0.0
Dental	90.7	90.7	0.0
OB-GYN	99.8	99.8	0.0
Primary Care	89.2	92.1	2.9
Psychiatry	99.6	99.8	0.2
Surgery	99.7	99.7	0.0

Note:

Desert: county-specialty-quarter with <1 active provider per 10,000 population. Difference = MD/DO rate minus all-clinicians rate (positive = including NPs reduces desert rate).

six specialties ($p > 0.10$ in all cases). Post-treatment dynamics vary modestly in magnitude but no specialty shows a statistically significant departure from zero. These results mirror the pooled event study in Figure 13 and are consistent with the specialty-level null reported in Table 4.

B.4 Alternative Clustering

Our baseline clusters standard errors at the state level (51 clusters). We verify robustness to alternative clustering levels: (a) state \times specialty (306 clusters) yields slightly smaller standard errors without changing the null conclusion; (b) two-way clustering by state and quarter yields virtually identical inference. With any clustering scheme, no specification achieves statistical significance.