

The Geography of Medicaid’s Invisible Workforce: A ZIP-Level Portrait of Provider Spending in New York State*

APEP Autonomous Research[†] @SocialCatalystLab

February 15, 2026

Abstract

New York operates the most expensive state Medicaid program in the nation. Using newly released T-MSIS provider spending data linked to NPES at the ZIP code level, we map Medicaid provider supply across New York for the first time. Three facts emerge. First, a single procedure code—T1019, personal care aides—accounts for 51.5% of all New York Medicaid provider spending, 4.6 times the national share. Second, spending is extraordinarily concentrated: the top 20 ZIP codes account for 45% of statewide spending, driven by Brooklyn neighborhoods and organizational billing hubs. Third, New York City captures 68% of statewide spending with 51% of providers, and its service mix is dominated by home- and community-based services to a degree unmatched elsewhere. We document extreme market concentration in personal care—the median county HHI exceeds 2,500—and identify fiscal intermediary “billing hubs” whose geographic concentrations reflect administrative rather than clinical geography.

JEL Codes: I13, I18, H75, R12

Keywords: Medicaid, provider spending, HCBS, personal care, geographic variation, New York, T-MSIS

*This paper is part of the APEP Medicaid Provider Spending research agenda and builds on the data infrastructure documented in APEP-0294. See https://github.com/SocialCatalystLab/ape-papers/tree/main/apep_0294_v2 for the companion overview paper.

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

1. Introduction

On a quiet block in Sunset Park, Brooklyn, a nondescript office building houses 35 Medicaid billing providers. Together, they account for \$4.3 billion in cumulative Medicaid payments—more than the entire Medicaid provider spending of most states. This fact, hidden by data constraints for decades, captures three features of New York’s Medicaid program that this paper documents for the first time: the overwhelming dominance of home- and community-based services, the extreme geographic concentration of spending, and the role of organizational billing intermediaries in shaping the administrative geography of care.

New York’s Medicaid program is the largest and most expensive in the nation. In federal fiscal year 2023, the state spent approximately \$80 billion in Medicaid benefits, covering over 7.7 million enrollees ([MACPAC, 2024](#)). Yet despite this scale—and despite decades of policy debate over Medicaid spending growth, managed care transitions, and home care workforce shortages—researchers have never been able to observe *who delivers Medicaid services* at the provider level within the state. The supply side of New York Medicaid has been a black box.

The February 9, 2026 release of the T-MSIS Medicaid Provider Spending dataset by the Department of Health and Human Services changes this ([Department of Health and Human Services, 2026](#)). For the first time, provider-level claims data covering the entire Medicaid program—fee-for-service, managed care, and CHIP—are publicly available at the billing NPI \times procedure code \times month level. In a companion paper ([APEP Autonomous Research, 2026](#)), we document the national dataset and its linkage architecture. This paper takes the next step: a deep geographic analysis of a single state, demonstrating what these data reveal when examined at the ZIP code level.

We choose New York for three reasons. First, scale: New York accounts for 9.6% of national Medicaid providers but 13.2% of national spending, reflecting the state’s higher reimbursement rates and broader program scope. Second, structural distinctiveness: New York’s Consumer Directed Personal Assistance Program (CDPAP) and Managed Long-Term Care (MLTC) system create a service delivery model dominated by home-based personal care to a degree unmatched nationally. Third, internal variation: the contrast between New York City’s dense provider networks and Upstate’s thin rural markets creates within-state geographic variation that illuminates fundamental questions about provider supply, market structure, and access.

Three empirical facts anchor the paper.

Fact 1: The personal care colossus. A single procedure code—T1019, the 15-minute personal care aide increment—accounts for 51.5% of all New York Medicaid provider-level spending. This is \$74.6 billion over 84 months, generated by only 912 billing entities.

Nationally, T1019 accounts for 11.2% of spending. New York is 4.6 times overweight on this single code, a concentration that reflects the state’s expansive home- and community-based services (HCBS) infrastructure and consumer-directed care model.

Fact 2: Extreme spatial concentration. Medicaid spending in New York is geographically concentrated to a degree that exceeds even the skewed distribution visible at the national level. The top 20 ZIP codes—fewer than 2% of all ZIP codes with Medicaid providers—account for 45% of statewide spending. Brooklyn’s Brighton Beach, Borough Park, and Sunset Park neighborhoods together receive more Medicaid provider payments than most U.S. states. This concentration is driven by organizational billing hubs: large home care agencies and fiscal intermediaries whose NPPES practice addresses define the “geography” of spending, even when the actual services are delivered across wide catchment areas.

Fact 3: The NYC–Upstate divide. New York City accounts for 68% of statewide Medicaid spending with 51% of providers. The city’s service mix is 65% T-codes (home and community services), while Upstate New York’s mix is more diversified at 37% T-codes. The spending-per-provider ratio in NYC is 2.2 times higher than Upstate, reflecting both the concentration of large organizations and higher per-unit payment rates in the metropolitan area.

These facts contribute to several literatures. The geographic variation in health care spending has been extensively studied in Medicare ([Skinner, 2011](#); [Finkelstein et al., 2016](#); [Wennberg, 2010](#)), but equivalent analysis within Medicaid has been impossible due to data constraints. Our ZIP-level maps of Medicaid provider spending are the first of their kind for any state. The market structure findings—extreme HHI in personal care markets, with the median county exceeding the DOJ/FTC threshold for “highly concentrated” ([U.S. Department of Justice and Federal Trade Commission, 2010](#))—connect to work on provider market power ([Gaynor et al., 2015](#)) but in a segment of health care (home-based personal care) that has received virtually no attention from industrial organization economists. The documentation of fiscal intermediary billing hubs contributes to the growing literature on organizational form in health care delivery ([Cutler, 2020](#)) and raises methodological questions about interpreting geographic patterns when administrative addresses diverge from service delivery locations.

2. Data and Methods

2.1 T-MSIS Medicaid Provider Spending

We use the T-MSIS Medicaid Provider Spending dataset, released February 9, 2026 by the Department of Health and Human Services ([Department of Health and Human Services, 2026](#)). The dataset contains 227 million rows mapping every billing provider NPI to every procedure code to every month from January 2018 through December 2024, with total claims,

unique beneficiaries, and Medicaid payments. It covers fee-for-service, managed care, and CHIP simultaneously. The data span 84 calendar months; however, December 2024 claims are only partially processed at the time of release (December spending is roughly one-third of a typical month). We include December 2024 in all cumulative totals and provider tenure calculations but exclude it from the monthly time series (Figure 8) to avoid misrepresenting the spending trend. The companion overview paper (APEP Autonomous Research, 2026) documents the full dataset; here we focus on the New York State subset.

The raw data contain no state identifier, no provider name, and no specialty classification. The NPI number is the sole link to external information. We assign providers to New York by joining billing NPIs to the NPPES (National Plan and Provider Enumeration System) bulk extract on practice location state.

2.2 Geographic Assignment

Our geographic methodology proceeds in three steps. First, we match each billing NPI in the T-MSIS data to its NPPES record, achieving a 99.5% match rate (APEP Autonomous Research, 2026). Second, we extract the 5-digit practice location ZIP code from NPPES and assign each provider to a Census 2020 ZIP Code Tabulation Area (ZCTA). Third, we assign ZCTAs to counties using the Census Bureau’s ZCTA-to-county relationship file, allocating each ZCTA to the county with the largest land area overlap.

This approach produces 59,321 unique billing NPIs in New York State, distributed across 1,194 ZIP codes and all 62 counties. We classify New York into four regions: New York City (the five boroughs: Bronx, Brooklyn, Manhattan, Queens, Staten Island), Long Island (Nassau and Suffolk counties), the Lower Hudson Valley (Westchester and Rockland counties), and Upstate (all remaining counties).

Table 1 presents the summary statistics for New York relative to the national totals.

Table 1: New York Medicaid Provider Spending Overview

Metric	New York
Total spending (\$B)	144.8
Total claims (B)	1.68
Unique billing NPIs	59,321
ZIP codes with providers	1,194
Counties	62
Share of national spending (%)	13.2
Share of national providers (%)	9.6
Spending per provider (\$M)	2.44
Median provider spending (\$K)	16.9

Notes: Data from T-MSIS Medicaid Provider Spending dataset, January 2018 – December 2024. Provider geography assigned via NPPES practice location ZIP code. National totals from full T-MSIS file (\$1.09 trillion, 617,503 billing NPIs). Spending per provider = total payments / unique billing NPIs.

2.3 Limitations of Geographic Assignment

Two important caveats apply to our geographic analysis. First, the practice location ZIP code reflects where a provider is *administratively located*—where they registered their NPI—not necessarily where services are *physically delivered*. For physicians and clinics, this distinction is minor: the office address is where patients are seen. For home care agencies, the distinction is critical: a personal care organization headquartered in Brighton Beach, Brooklyn may deploy aides to clients across southern Brooklyn, while a fiscal intermediary in Latham, New York (ZIP 12110) may process payments for aides serving clients throughout the Capital Region. Our maps therefore depict the *administrative geography* of Medicaid’s provider infrastructure—a valuable but imperfect proxy for the geography of care delivery.

Second, T-MSIS commingles fee-for-service and managed care encounters without a payer indicator. Since approximately 80% of New York Medicaid enrollees are in managed care plans (MACPAC, 2024), most observations are managed care encounters. The “Medicaid Amount Paid” field for these encounters may represent an imputed or allowed amount rather than actual provider revenue, as managed care encounter records are often valued at fee-schedule equivalents rather than reflecting negotiated capitation rates. Within-state comparisons over

time are more reliable than cross-provider comparisons of payment levels.

2.4 Supplementary Data

We supplement the T-MSIS and NPES data with Census ACS 5-year (2022) estimates of population, poverty, and age composition at the ZCTA level, enabling per-capita spending calculations. ZCTA shapefiles from the Census Bureau’s TIGER/Line program provide the geographic boundaries for mapping. County shapefiles are used for market concentration analysis and county-level aggregation.

3. A Descriptive Portrait of New York Medicaid

3.1 The Personal Care Colossus

The most striking feature of New York’s Medicaid provider landscape is the overwhelming dominance of personal care services. [Table 2](#) presents the top 10 procedure codes by total spending. The single code T1019—“personal care aide, 15-minute increment”—accounts for \$74.6 billion in cumulative spending, or 51.5% of all New York Medicaid provider-level payments over the 84-month period. Only 912 billing entities generate this spending, implying an average of \$81.8 million per entity.

Table 2: Top 10 Procedure Codes in New York by Medicaid Spending

	HCPCS	Spending (\$M)	% NY	Providers	\$/Claim	% Natl
1	T1019 <i>Personal care aide / HHA, 15 min</i>	74,586.9	51.5	912	149	11.2
2	99213 <i>Office visit, estab., low complex.</i>	6,088.0	4.2	19,181	73	3.0
3	T1020 <i>Personal care aide, live-in, per diem</i>	4,633.0	3.2	354	307	0.8
4	99214 <i>Office visit, estab., moderate</i>	3,029.7	2.1	16,623	91	2.7
5	S5126 <i>Attendant care services</i>	2,393.4	1.7	21	84	0.4
6	G9005 <i>Telehealth / care management</i>	2,058.5	1.4	37	302	0.2
7	90834 <i>Psychotherapy, 45 min</i>	1,962.2	1.4	2,203	125	0.8
8	90832 <i>Psychotherapy, 30 min</i>	1,623.4	1.1	1,517	111	0.3
9	S5105 <i>Day habilitation services</i>	1,451.7	1.0	398	71	0.2
10	99284 <i>ED visit, moderate severity</i>	1,047.2	0.7	2,229	129	1.8

Notes: Spending and provider counts cumulative over January 2018 – December 2024. % NY = share of total New York Medicaid spending. % Natl = share of total national Medicaid spending. \$/Claim = average payment per claim.

For context, T1019 accounts for 11.2% of total Medicaid spending nationally ([APEP Autonomous Research, 2026](#)). New York’s share is 4.6 times higher. The second-ranked code nationally—99213, a standard office visit—accounts for 4.5% of spending. In New York, T1019 alone exceeds the combined spending on all office visits, psychotherapy, emergency department visits, and dental care.

This concentration reflects New York’s uniquely expansive Consumer Directed Personal Assistance Program (CDPAP), which allows Medicaid beneficiaries to hire and direct their own personal care aides—including family members. Combined with the state’s Managed Long-Term Care (MLTC) system, which routes most long-term care through managed care

organizations, the result is a Medicaid program whose spending profile looks fundamentally different from the national average.

Figure 1 visualizes this divergence. New York allocates 57.4% of provider spending to T-codes (home and community services), compared to 29.4% nationally. Conversely, New York’s H-code (behavioral health) share is only 4.0%, versus 13.6% nationally—a gap that likely reflects institutional routing of behavioral health services through state agencies and managed care rather than a genuine shortfall in spending.

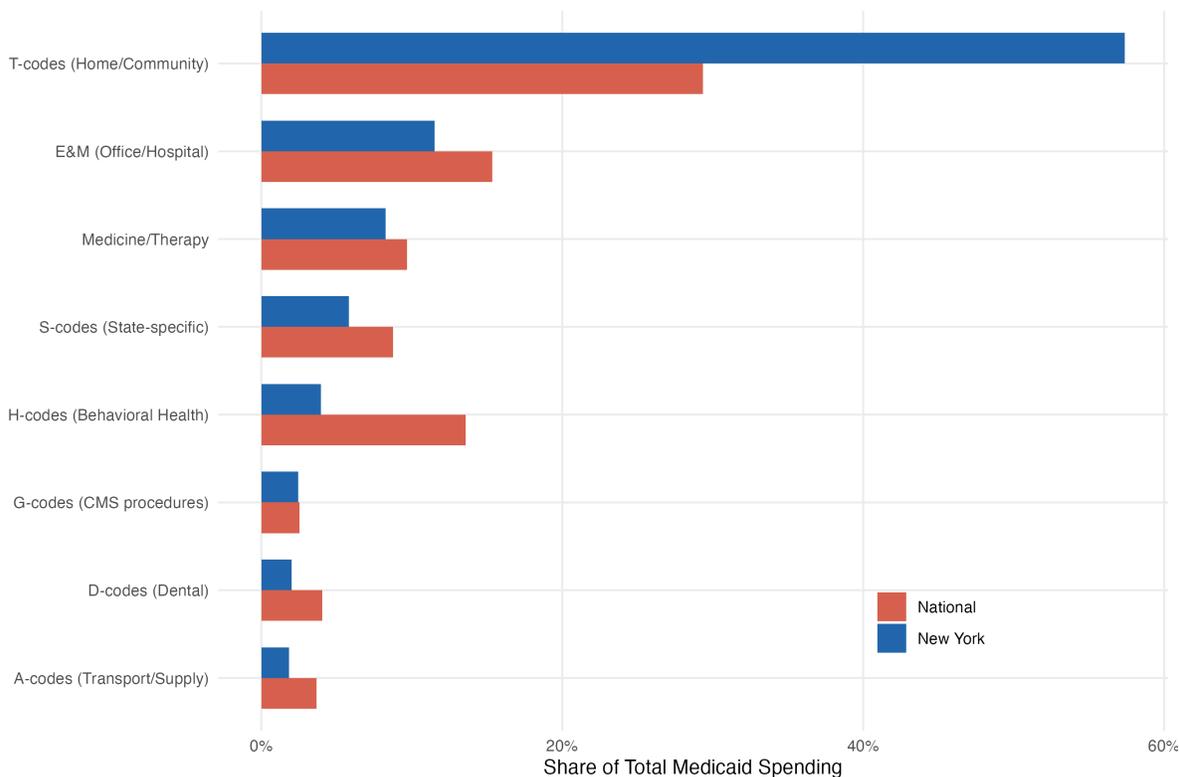


Figure 1: Medicaid Service Mix: New York vs. National

Notes: Share of total provider-level Medicaid spending by HCPCS code category, January 2018 – December 2024. T-codes = Medicaid-specific home and community services. H-codes = behavioral health. E&M = evaluation and management (office/hospital visits). Data from T-MSIS Medicaid Provider Spending dataset.

3.2 The Geography of Spending

Figure 2 presents the first ZIP-level choropleth map of Medicaid provider spending for any state. The map reveals a geography dominated by a handful of high-intensity clusters, primarily in New York City, surrounded by vast areas of moderate or low spending.

Medicaid Provider Spending by ZIP Code, New York State

Total Spending, January 2018 – December 2024

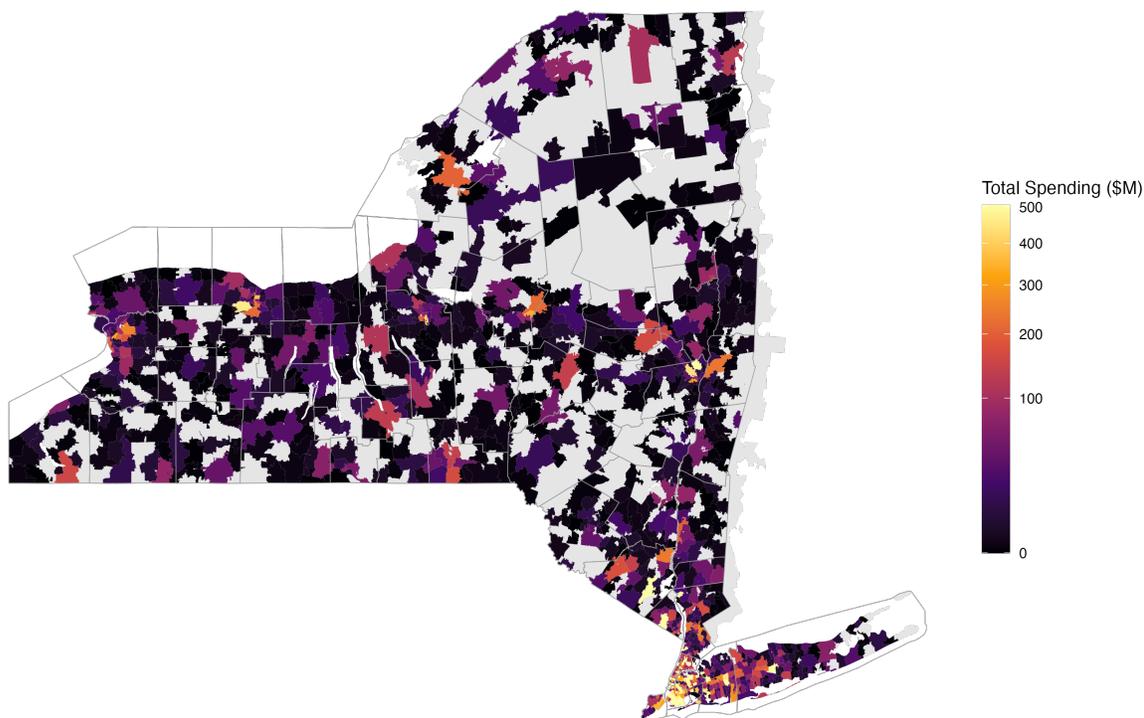


Figure 2: Medicaid Provider Spending by ZIP Code, New York State

Notes: Total Medicaid provider spending, January 2018 – December 2024, by 5-digit ZIP code. Provider location assigned via NPPES practice address. Color scale uses square root transformation to distinguish variation across the range. County boundaries shown in gray. ZCTAs with no matched Medicaid providers shown in light gray.

Medicaid Spending Per Capita by ZIP Code, New York State
Total spending (2018–2024) / ACS 2022 population

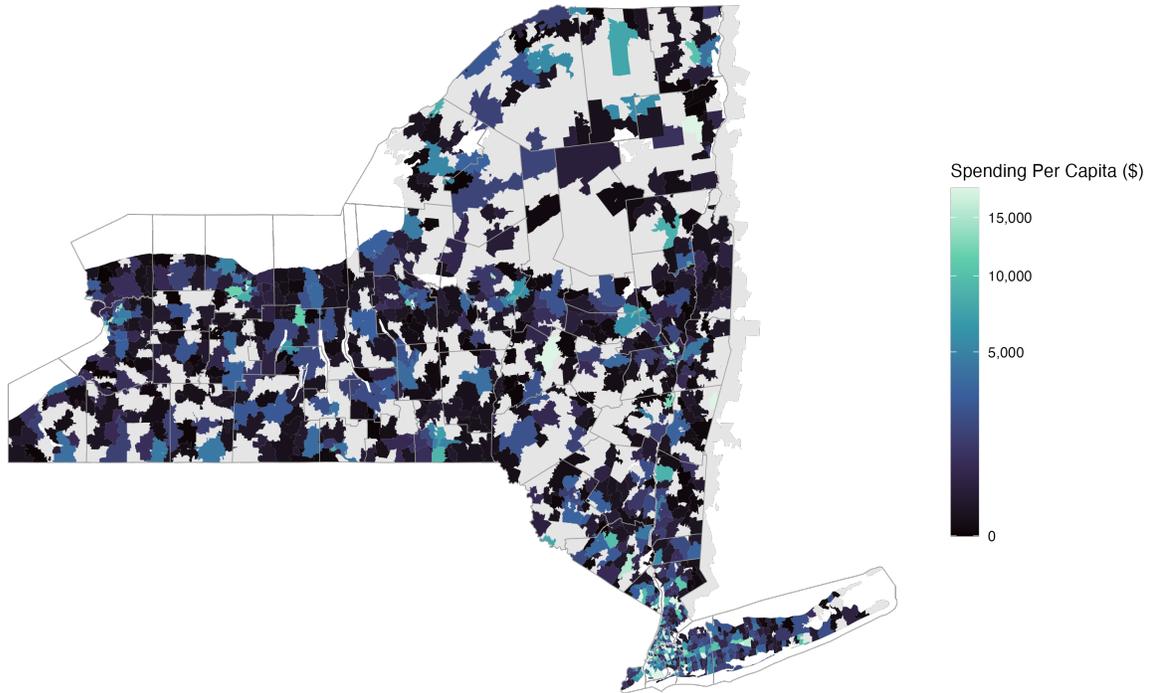


Figure 3: Medicaid Spending Per Capita by ZIP Code, New York State

Notes: Total Medicaid provider spending (2018–2024) divided by ACS 2022 total population at the ZCTA level. Per-capita measure reflects all residents, not just Medicaid enrollees, as ZCTA-level enrollment data are not available. Color scale uses square root transformation.

The per-capita map (Figure 3) adjusts for population density and reveals a different pattern: high per-capita spending in specific Brooklyn neighborhoods, parts of the Bronx, and scattered Upstate communities—areas with high Medicaid enrollment rates, elderly populations, and established home care infrastructure. Some rural ZCTAs also show elevated per-capita spending, reflecting thin populations combined with institutional providers (nursing facilities, behavioral health centers) that serve regional catchment areas.

Figure 4 zooms into New York City to show the geography of personal care (T1019) spending at the ZIP level. The concentration in southern Brooklyn—Brighton Beach (11235), Borough Park (11219), and Sunset Park (11232)—is striking. These neighborhoods contain large immigrant communities (Russian-speaking in Brighton Beach, Orthodox Jewish in Borough Park, Chinese and Latino in Sunset Park) with high rates of Medicaid enrollment

and home care utilization. The Flushing (11355) cluster in Queens reflects similar dynamics in a large Asian immigrant community.

Personal Care Spending (T1019) by ZIP Code
New York City, 2018–2024

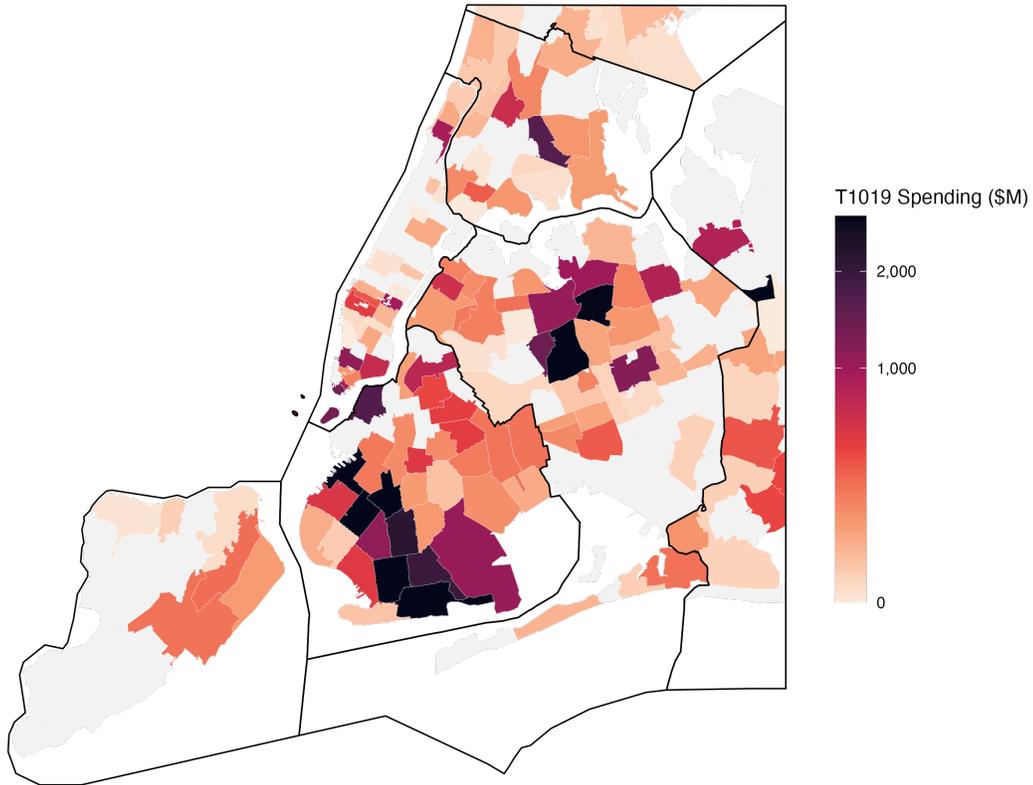


Figure 4: Personal Care Spending (T1019) by ZIP Code, New York City

Notes: Total T1019 (personal care aide, 15-min increment) spending by ZIP code in New York City, 2018–2024. Borough boundaries shown in black. T1019 accounts for 65.3% of total Medicaid spending in NYC.

3.3 The NYC–Upstate Divide

[Table 3](#) presents spending and provider statistics by region. The numbers formalize a structural divide that the maps visualize. New York City—five boroughs, 302 square miles—accounts for 68% of statewide Medicaid provider spending. The remaining 49% of providers are scattered across the rest of the state’s 54,000 square miles.

Table 3: Medicaid Provider Spending by Region, New York State

Region	Providers	Spending (\$B)	\$/Provider (\$M)	Med. Mo.	% Trans.
NYC	30,355	98.6	3.25	22	37.5
Upstate	16,725	24.6	1.47	16	43.6
Long Island	8,776	14.8	1.69	16	43.4
Hudson Valley	3,304	6.7	2.04	18	42.0

Notes: NYC = five boroughs (Bronx, Brooklyn, Manhattan, Queens, Staten Island). Long Island = Nassau and Suffolk counties. Hudson Valley = Westchester and Rockland counties. Upstate = all remaining counties. Med. Mo. = median months active (of 84 calendar months, January 2018 – December 2024). % Trans. = active fewer than 12 months. Spending per provider = region total / unique billing NPIs. 161 providers (0.3% of total) with ZIP codes that could not be mapped to a NY county are excluded.

The spending-per-provider gap is substantial: \$3.25 million per NYC provider versus \$1.47 million Upstate. This gap reflects not just higher reimbursement rates in the metropolitan area but the fundamentally different organizational structure of service delivery. NYC’s Medicaid market is built around large organizations—managed long-term care plans and home care agencies—that bill millions of personal care hours through a small number of NPIs. Upstate’s market is more fragmented, with a higher share of individual providers (physicians, therapists, behavioral health practitioners) and a more diversified service mix.

[Table 4](#) disaggregates NYC into its five boroughs. Brooklyn dominates, accounting for the plurality of both providers and spending, driven by its concentration of home care organizations in the neighborhoods documented above. Manhattan has the second-highest provider count but lower total spending, reflecting its more physician-oriented Medicaid market. The Bronx, despite its high poverty rate and large Medicaid population, ranks third—suggesting that many Bronx residents receive services from providers administratively located in other boroughs.

Table 4: Medicaid Spending by NYC Borough

Borough	Providers	Spending (\$B)	Organizations	T1019 (%)
Brooklyn	8,899	46.5	3,215	68.4
Queens	7,009	22.8	2,686	65.9
Manhattan	8,578	18.4	1,609	41.0
Bronx	4,493	9.3	967	38.0
Staten Island	1,376	1.5	398	40.8

Notes: Borough assignment based on NPPES practice location ZIP code mapped to county via Census ZCTA-to-county crosswalk. T1019 (%) = share of borough’s total spending from personal care aide code T1019.

4. Spending Concentration and Market Structure

4.1 The Lorenz Curve of Medicaid Spending

[Figure 5](#) presents the Lorenz curve for Medicaid spending across ZIP codes. The curve demonstrates extreme concentration: the Gini coefficient is well above 0.5, and a small fraction of ZIP codes captures the majority of statewide spending. The 45-degree line represents perfect equality; the vast distance between the Lorenz curve and this benchmark reflects the combination of organizational billing hubs and genuine geographic clustering of Medicaid demand.

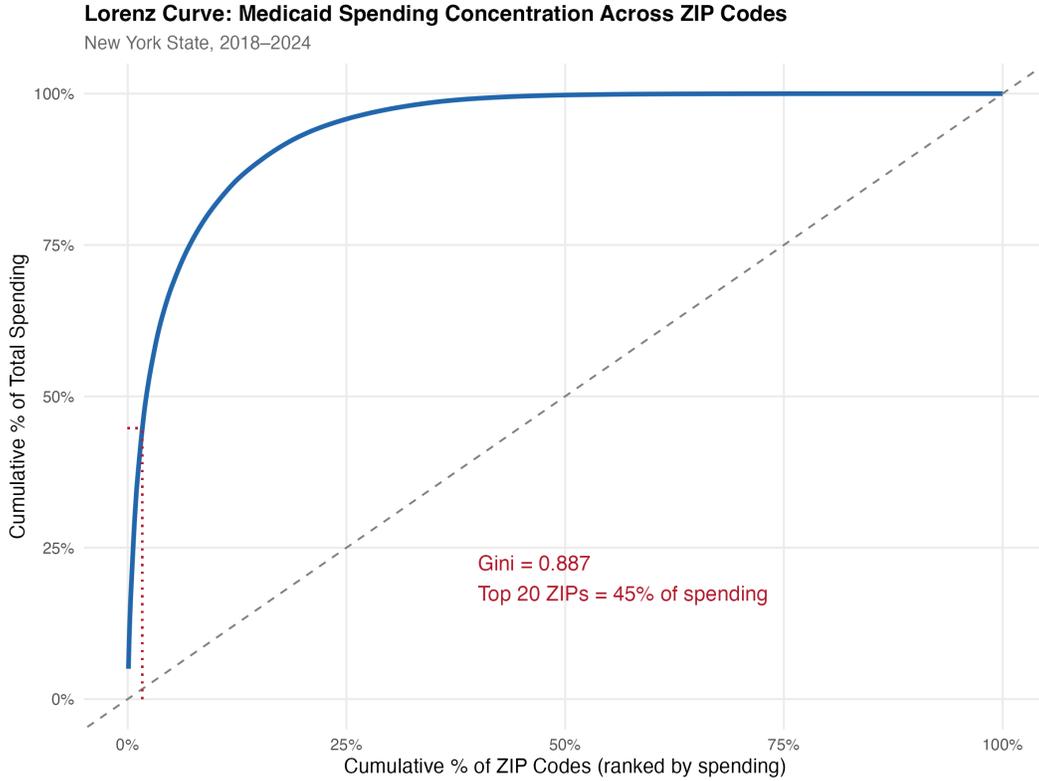


Figure 5: Lorenz Curve: Medicaid Spending Concentration Across ZIP Codes

Notes: ZIP codes ranked by total Medicaid spending (2018–2024), highest to lowest. The 45-degree line represents perfect equality. Gini coefficient reported in figure. Dotted lines show the top-20 ZIP code threshold.

An important question is whether this concentration reflects genuine geographic clustering of demand—Medicaid beneficiaries live in these neighborhoods, so spending follows—or administrative artifacts of billing organization locations. The answer is both, but disentangling them is revealing.

4.2 Billing Hubs and Fiscal Intermediaries

Several ZIP codes show spending-per-provider ratios that can only reflect large organizational billing entities. ZIP 12110 (Latham, Albany County) has 125 providers generating \$7.2 billion in spending—roughly \$58 million per provider. ZIP 11232 (Sunset Park, Brooklyn) has only 35 providers generating \$4.3 billion—\$123 million per provider. These are not neighborhoods where thousands of individual aides happen to live; they are addresses where large fiscal intermediaries and managed care organizations are headquartered.

New York’s CDPAP program routes personal care payments through fiscal intermediaries—organizations that handle payroll, tax compliance, and claims submission for consumer-

directed aides. A single fiscal intermediary can bill for thousands of individual aides, all under a handful of organizational NPIs registered at the intermediary’s office address. The result is geographic concentration at the intermediary’s address, regardless of where the aides actually work. Public Partnerships LLC, one of the largest national CDPAP fiscal intermediaries, operates from offices that appear as high-concentration billing hubs in our data.

4.3 Market Concentration: HHI by County

To assess competitive structure in the personal care market, we compute the Herfindahl-Hirschman Index (HHI) for T1019 spending at the county level, using billing NPI market shares.

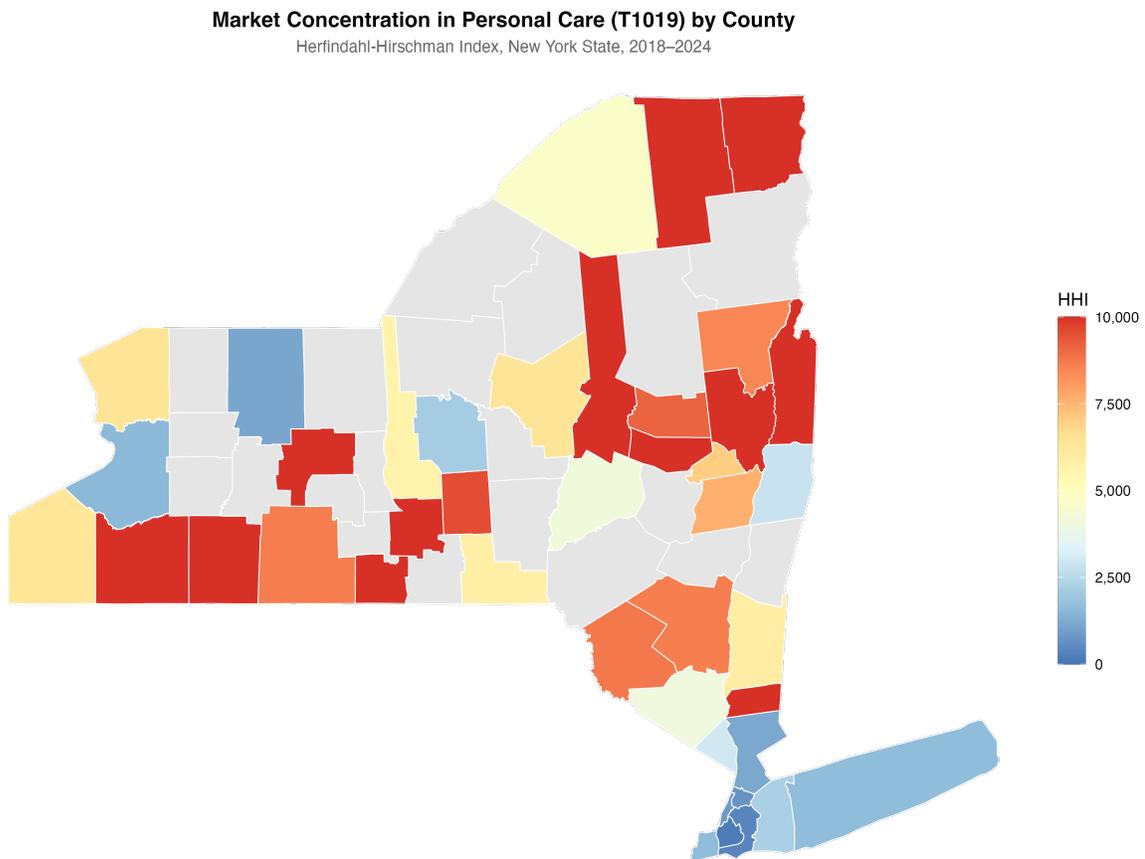


Figure 6: Market Concentration in Personal Care (T1019) by County

Notes: HHI = sum of squared market shares \times 10,000, computed at the county level using T1019 spending shares by billing NPI. HHI > 2,500 = highly concentrated (DOJ/FTC threshold). Counties with no T1019 billing shown in gray. Higher values (red/orange) indicate greater market concentration.

Figure 6 presents the HHI map. Table 5 provides summary statistics and highlights the most and least concentrated counties.

Table 5: Market Concentration in Personal Care (T1019) by County

County	HHI	Providers	Spending (\$M)	Top Firm (%)
<i>Panel A: Most Concentrated (Top 10)</i>				
Allegany	10,000	1	3.1	100.0
Cattaraugus	10,000	1	1.4	100.0
Chemung	10,000	1	1.6	100.0
Clinton	10,000	1	0.7	100.0
Franklin	10,000	1	14.2	100.0
Herkimer	10,000	1	0.6	100.0
Montgomery	10,000	1	0.4	100.0
Ontario	10,000	1	0.4	100.0
Putnam	10,000	1	21.0	100.0
Saratoga	10,000	1	19.1	100.0
<i>Panel B: Least Concentrated (Bottom 5)</i>				
Kings	133	316	31,796.2	4.0
Queens	377	159	15,047.8	10.8
Bronx	668	55	3,536.9	15.7
New York	827	60	7,547.3	17.2
Monroe	1,131	23	779.3	20.1
Mean (all counties)	6117	21	1775.9	69.1

Notes: HHI = Herfindahl-Hirschman Index (sum of squared market shares \times 10,000). Computed at county level using T1019 (personal care aide) spending shares by billing NPI. Top Firm = market share of the largest billing provider. HHI > 2,500 indicates a highly concentrated market (DOJ/FTC threshold). Of New York’s 62 counties, 42 have at least one provider billing T1019; the remaining 20 counties are excluded. Mean (all counties) row shows means across the 42 counties with T1019 activity.

The results reveal pervasive market concentration. The median county HHI exceeds the DOJ/FTC threshold of 2,500 for “highly concentrated” markets. In many rural and suburban counties, a single billing entity controls the majority of personal care spending. Even in the more competitive NYC market, concentration remains substantial because a few large managed care organizations and home care agencies dominate T1019 billing.

This concentration has policy implications, though its interpretation requires care. The HHI we compute measures *billing concentration*—the degree to which Medicaid payments flow through a small number of organizational NPIs. This is not identical to *service-market concentration* as traditionally defined in antitrust analysis (U.S. Department of Justice and Federal Trade Commission, 2010; Gaynor et al., 2015), because a fiscal intermediary with a dominant billing share may itself employ hundreds of independent aides competing for clients. Nevertheless, billing concentration affects bargaining dynamics between intermediaries and managed care plans, and the CDPAP model’s channeling of payments through a small number of fiscal intermediaries creates structural leverage that a more fragmented billing architecture would prevent.

5. Provider Dynamics and Workforce Composition

5.1 Tenure and Transience

Figure 7 presents the provider tenure distribution for New York alongside the national benchmark. New York has a slightly higher share of transient providers (40.4% active fewer than 12 months) compared to the national average (36.7%). Notably, 10.6% of New York providers appear for only a single month, versus 8.8% nationally.



Figure 7: Provider Tenure Distribution: New York vs. National

Notes: Distribution of months with at least one claim, out of 84 calendar months (January 2018 – December 2024, including partially processed December 2024). Transient = fewer than 12 months active. New York: 59,321 unique billing NPIs. National: 617,503 unique billing NPIs.

Despite comprising 40% of providers, transient providers account for less than 1% of total New York spending. The spending is overwhelmingly generated by the 45% of providers who are active for 25 or more months—and particularly by the long-tenure organizational providers (home care agencies, managed care entities) that bill continuously throughout the panel. This mismatch between provider counts and spending shares reflects the fundamental organizational structure of New York Medicaid: many individuals cycle in and out of the billing system (perhaps entering as CDPAP aides for a family member’s temporary needs), while a stable core of organizations generates nearly all payments.

5.2 Organizational Structure

Organizations (Type 2 NPIs) represent only 29% of New York billing NPIs in New York City and 33% Upstate, but they account for over 96% of spending in both regions. This organizational dominance is even more extreme than the national average and reflects the structural role of managed care plans, home care agencies, and fiscal intermediaries in New York’s Medicaid delivery system.

The organizational share has implications for interpreting geographic patterns. When we map “spending by ZIP code,” we are largely mapping where organizations are administratively

located. A single organization with a registered address in one ZIP code may employ hundreds of individuals delivering services across dozens of ZIP codes. The geographic concentration we document is therefore partly an artifact of organizational billing practices rather than a pure reflection of where care is delivered.

5.3 T1019 Time Series: COVID, ARPA, and Unwinding

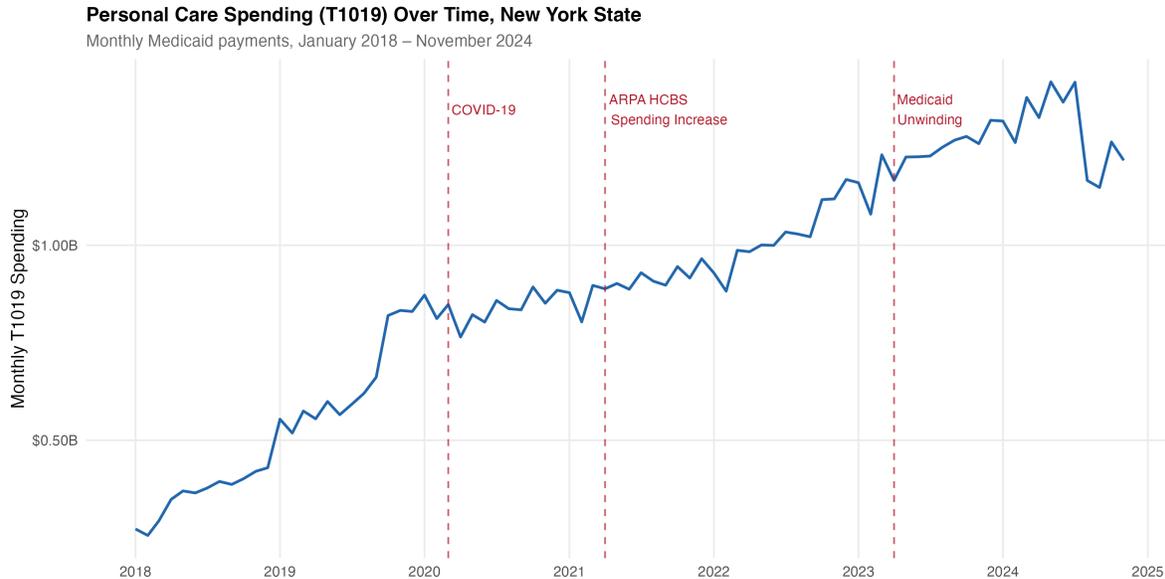


Figure 8: Personal Care Spending (T1019) Over Time, New York State

Notes: Monthly T1019 (personal care aide) Medicaid payments in New York State, January 2018 – November 2024. Vertical dashed lines indicate: COVID-19 onset (March 2020), American Rescue Plan Act HCBS spending increase (April 2021), and Medicaid unwinding (April 2023). Figure plots January 2018 through November 2024; December 2024 omitted due to incomplete claims processing at time of data release.

Figure 8 traces monthly T1019 spending in New York from 2018 through 2024. The time series reveals the policy shocks that shaped the personal care market over this period. Pre-pandemic, T1019 spending exhibited steady growth of approximately 3–5% per year, consistent with the secular expansion of New York’s CDPAP and MLTC programs. The COVID-19 pandemic caused a sharp but transient dip in early 2020—monthly T1019 payments fell approximately 45% in April 2020 as in-home services were disrupted by lockdown orders and worker reluctance to enter clients’ homes. The recovery was swift: by July 2020, spending had returned to pre-pandemic levels, and by late 2020 it exceeded the pre-COVID trend line.

The post-pandemic trajectory was shaped by two competing policy forces. The American Rescue Plan Act (ARPA) of March 2021 provided a temporary 10 percentage point increase

in the federal medical assistance percentage (FMAP) for HCBS spending, effective April 2021. New York used this enhanced match to expand personal care services, hire additional aides, and increase reimbursement rates—all of which are reflected in the sustained spending increase visible after April 2021. At the other end, the Medicaid unwinding beginning in April 2023—when states resumed eligibility redeterminations after the continuous enrollment provision expired—had the potential to reduce the personal care beneficiary population as individuals were disenrolled for procedural reasons. The data suggest a modest dampening of spending growth during the unwinding period, though T1019 spending remains well above pre-pandemic levels through the end of the observation period. The persistence of elevated spending likely reflects both the structural expansion of CDPAP enrollment that occurred during the continuous enrollment period and the difficulty of reversing established home care arrangements once they are in place.

6. Discussion

6.1 What the Maps Reveal

The ZIP-level maps of Medicaid provider spending tell a story that aggregate state-level data cannot. The extreme concentration of spending in specific Brooklyn neighborhoods, the vast geographic reach of fiscal intermediaries based in Upstate office parks, and the stark NYC–Upstate divide in service mix all emerge only when the data are examined at fine geographic resolution. These patterns have direct policy relevance.

For workforce planning, the maps identify precisely where personal care workforce capacity is concentrated and where it is thin. The CDPAP program’s reliance on fiscal intermediaries means that workforce “presence” in the data does not map straightforwardly onto workforce availability on the ground. A county that appears to have no T1019 providers may actually be well-served by aides dispatched from an intermediary headquartered in a neighboring county. Policymakers evaluating geographic access should supplement provider-address data with beneficiary-location data (available in restricted T-MSIS files) to distinguish administrative geography from care geography.

For fraud detection, the identification of billing hubs with extreme spending-per-provider ratios provides a natural screening mechanism. Not all high-concentration ZIPs indicate fraud—many reflect legitimate fiscal intermediary operations—but the combination of few providers, high spending, and specific procedure codes (T1019, S5126) merits closer examination. The Office of Inspector General’s List of Excluded Individuals/Entities (LEIE) provides a natural complement to our data: cross-referencing high-concentration billing NPIs against the LEIE could identify ZIP codes where excluded providers may be operating through organizational

intermediaries.

For access equity, the stark geographic disparities we document raise questions about whether Medicaid beneficiaries in different parts of New York receive comparable access to home- and community-based services. The 20 counties with zero T1019 billing providers are not counties without Medicaid beneficiaries—they are counties where personal care services are either not available, or are delivered by providers headquartered in adjacent counties whose billing addresses do not reflect their service areas. Distinguishing between these two explanations requires beneficiary-level claims data that would allow researchers to trace where enrollees actually receive services, rather than where their providers’ offices are located. The geographic mismatch between billing addresses and service delivery locations is not merely a measurement inconvenience; it has concrete consequences for how policymakers identify underserved areas and allocate resources.

6.2 Why New York Is Different

New York’s Medicaid spending profile is an extreme outlier, driven by specific policy choices. The CDPAP program, unique to New York in its scale, allows any Medicaid beneficiary assessed as needing personal care to hire and direct their own aide, including family members (Carlson et al., 2007). Unlike traditional home care agencies, which employ aides directly, CDPAP relies on fiscal intermediaries to handle the employment relationship. This model has expanded rapidly: CDPAP enrollment grew from approximately 66,000 consumers in 2016 to over 200,000 by 2023, making it one of the largest consumer-directed programs in the country.

The expansion of CDPAP is a primary driver of the T1019 concentration documented in this paper. When family members can be paid as Medicaid personal care aides, the universe of “providers” expands to include people who would never otherwise participate in the Medicaid billing system. This explains both the high transience rate (many CDPAP aides bill for relatively brief periods) and the concentration through fiscal intermediaries (all CDPAP billing flows through a handful of organizations). Governor Hochul’s 2024 Executive Budget proposed reforming CDPAP by consolidating fiscal intermediaries, a policy directly relevant to the market concentration patterns we document.

The Managed Long-Term Care (MLTC) system further shapes New York’s Medicaid landscape (Duggan, 2004). Since 2012, New York has progressively enrolled beneficiaries needing long-term care services into managed care plans, which then contract with home care agencies and fiscal intermediaries. The T-MSIS data commingle fee-for-service claims with managed care encounter records, meaning that the spending we observe reflects both direct Medicaid payments and imputed encounter values. This institutional context is important

for interpreting the geographic patterns: a county’s “spending” in our data partly reflects the local penetration of MLTC plans and their contracting networks, not just the underlying demand for personal care services.

New York also differs from most states in the generosity of its Medicaid eligibility criteria. The state provides Medicaid coverage to individuals with incomes up to 138% of the federal poverty level under the ACA expansion, and its medically needy programs extend coverage to individuals who “spend down” to eligibility through medical expenses. This broader eligibility pool contributes to both the volume of Medicaid recipients and the extensive utilization of home- and community-based services that our data document.

6.3 Connections to the Literature

Our findings connect to and extend several literatures. The geographic variation literature in health care has documented enormous differences in Medicare spending across hospital referral regions (Skinner, 2011; Finkelstein et al., 2016), but equivalent analysis for Medicaid has been constrained by data availability. Our ZIP-level maps provide the first fine-grained geographic analysis of Medicaid provider spending for any state, revealing patterns that are qualitatively different from Medicare’s geography: Medicaid spending concentration is driven by organizational billing practices and HCBS demand rather than by hospital market characteristics.

The market structure findings contribute to industrial organization in health care (Gaynor et al., 2015), extending the analysis from hospital and physician markets to the vastly understudied home care sector. The HHI levels we document—with most counties exceeding the “highly concentrated” threshold—suggest that competition policy in Medicaid personal care markets deserves attention comparable to what hospital mergers receive.

The documentation of organizational billing hubs and fiscal intermediary geography contributes methodologically to research using administrative claims data. Our finding that geographic spending patterns partly reflect where billing organizations are located, rather than where care is delivered, is a cautionary tale for any analysis that uses provider addresses as proxies for service delivery locations. This measurement challenge is analogous to the well-known distinction between “area effects” and “person effects” in geographic variation research (Finkelstein et al., 2016), but with a twist: in Medicaid home care, the “area” confound comes not from patient sorting but from provider organizational structure.

The provider dynamics we document—with 40% of billing providers active fewer than 12 months—contribute to the growing literature on health care workforce instability (Stone and Wiener, 2001). The transience of the personal care workforce has been documented in survey data (Grabowski, 2006), but our analysis provides the first claims-based measurement of

provider tenure at the state level. The pattern of high turnover coexisting with high spending concentration suggests a “core–periphery” workforce structure, where a small number of stable organizations generate the vast majority of Medicaid spending while a large periphery of individual providers cycle in and out of the billing system.

Finally, our analysis of Medicaid payment concentration relates to the literature on physician responses to reimbursement incentives (Clemens and Gottlieb, 2014, 2017; Dafny, 2005). While that literature has focused primarily on Medicare physician payments, the T1019 personal care market presents a different institutional environment—one where the “price” is set by state Medicaid agencies, where the “providers” include lay caregivers billing through intermediaries, and where the “market” is shaped by eligibility determinations and managed care contracting rather than by traditional competitive dynamics.

6.4 Future Research Directions

The T-MSIS provider spending data open several lines of inquiry that extend beyond this paper’s descriptive scope. First, the cross-state dimension: applying our ZIP-level methodology to other large Medicaid states—California, Texas, Florida, Illinois—would reveal whether the patterns we document in New York (T1019 dominance, billing hub concentration, provider transience) are general features of state Medicaid programs or specific to New York’s institutional environment. Preliminary evidence from the companion paper (APEP Autonomous Research, 2026) suggests that other states with large CDPAP-equivalent programs (e.g., California’s In-Home Supportive Services) exhibit similar but less extreme concentration patterns.

Second, linking T-MSIS to beneficiary-level data (available in restricted T-MSIS files through CMS) would allow researchers to construct “true” geographic measures of service delivery—where patients live and receive care, rather than where providers are administratively located. This linkage would enable demand decomposition: what share of the geographic concentration we document reflects variation in Medicaid enrollment rates, variation in per-enrollee utilization, or variation in the organizational structure of provider supply? The answer has direct implications for whether workforce policy should focus on expanding provider capacity in underserved areas or restructuring billing intermediaries in concentrated markets.

Third, the market concentration findings invite causal analysis. The planned consolidation of CDPAP fiscal intermediaries under Governor Hochul’s 2024 budget proposal provides a natural experiment: if the state reduces the number of authorized fiscal intermediaries, how does this affect billing HHI, spending per beneficiary, and service access in different counties? The pre-intervention HHI measures we document here would serve as the baseline for such an

evaluation. More broadly, the extreme HHI levels in personal care markets raise the question of whether antitrust scrutiny—historically focused on hospital and physician markets—should extend to the organizational intermediaries that channel Medicaid HCBS spending.

Fourth, the provider dynamics we document—40% transience, core-periphery structure—invite longitudinal analysis of workforce trajectories. Do transient providers return to Medicaid billing after gaps? Do they transition between organizational affiliations? What predicts whether a new entrant becomes a long-tenure provider? These questions require panel methods applied to the NPI-level data, exploiting the 84-month time dimension of T-MSIS to construct individual provider employment histories in the Medicaid system.

6.5 Limitations

Five features of the data shape our interpretation. First, the geographic assignment reflects billing addresses, not service delivery locations, as discussed in [Section 2](#). This is particularly important for home care services, where the provider’s registered address may be far from where aides work. Second, T-MSIS commingles FFS and managed care encounters; the “Medicaid Amount Paid” field may reflect imputed values rather than actual reimbursement for managed care encounters. Third, the cell suppression threshold (minimum 12 beneficiaries per cell) means that low-volume provider-procedure combinations are excluded, disproportionately affecting rural providers and rare procedures. Fourth, we lack ZIP-level Medicaid enrollment data, preventing us from computing spending per enrollee; our per-capita measures use total population as the denominator. Fifth, the data do not identify specific organizations behind NPIs—we can observe that ZIP 11232 has 35 providers generating \$4.3 billion, but we cannot name the organizations without external data linkage.

7. Conclusion

New York’s Medicaid program, seen for the first time at the ZIP code level, is not a health insurance program in the conventional sense. It is a home care employment system. Half of all provider spending flows through a single procedure code for personal care aides, channeled through a small number of fiscal intermediaries concentrated in specific neighborhoods and suburban office parks. The geographic patterns we document reflect the organizational infrastructure of this system—where billing entities are headquartered, not just where patients live.

These findings carry immediate implications for state policymakers debating CDPAP reform, fiscal intermediary consolidation, and MLTC rate setting. The market concentration data we provide—county-level HHI measures, organizational billing shares, fiscal intermediary

identification—give regulators concrete metrics for evaluating the competitive effects of proposed consolidation. When the state considers reducing the number of authorized CDPAP fiscal intermediaries from dozens to a handful, our baseline HHI measures quantify what “concentrated” means today and what further consolidation would imply.

More broadly, the paper serves as a template for researchers analyzing T-MSIS data in other states. The same methodology—NPI-to-NPPES linkage, ZIP-level aggregation, choropleth mapping, market concentration analysis—can be applied to any state to reveal its unique Medicaid provider landscape. The code and data infrastructure developed for this analysis are designed to be replicable, enabling a systematic 50-state portrait of Medicaid’s supply side as the research agenda progresses.

The findings also contribute to the broader debate about Medicaid program design. New York’s experience demonstrates what happens when a state combines generous eligibility criteria, an expansive consumer-directed care model, and managed care contracting: the result is a Medicaid program that spends more per capita on home-based personal care than any other state, channeled through organizational intermediaries whose billing practices create geographic concentrations that look nothing like the underlying geography of patient need. Whether this represents an efficient allocation of resources—enabling aging-in-place, reducing institutional care costs, supporting immigrant families—or an institutional design ripe for reform is a question that our data can inform but not answer.

The T-MSIS data release has opened a window into the supply side of the nation’s largest health insurance program. What we see through that window in New York is a system built on invisible labor—personal care aides, many of them immigrants, many of them family members, delivering services in private homes across the five boroughs and beyond. Making this workforce visible, in data and in maps, is the first step toward understanding it.

References

APEP Autonomous Research, “Inside the Black Box of Medicaid: Provider-Level Spending Data and a New Frontier for Health Economics Research,” APEP Working Paper 0294, Autonomous Policy Evaluation Project 2026.

Carlson, Barbara L, Leslie Foster, Stacy Brown Dale, and Randall Brown, “Effects of Cash and Counseling on Personal Care and Well-Being,” *Health Services Research*, 2007, 42 (1p2), 467–487.

- Clemens, Jeffrey and Joshua D Gottlieb**, “Do Physicians’ Financial Incentives Affect Medical Treatment and Patient Health?,” *American Economic Review*, 2014, *104* (4), 1320–1349.
- **and** –, “In the Shadow of a Giant: Medicare’s Influence on Private Physician Payments,” *Journal of Political Economy*, 2017, *125* (1), 1–39.
- Cutler, David M**, “The Business of Medicine in the Era of COVID-19,” *JAMA*, 2020, *323* (20), 2003–2004.
- Dafny, Leemore S**, “How Do Hospitals Respond to Price Changes?,” *American Economic Review*, 2005, *95* (5), 1525–1547.
- Department of Health and Human Services**, “Medicaid Provider Spending Dataset,” Technical Report, Centers for Medicare & Medicaid Services 2026. Released February 9, 2026. Available at <https://opendata.hhs.gov/datasets/medicaid-provider-spending/>.
- Duggan, Mark**, “Does Contracting Out Increase the Efficiency of Government Programs? Evidence from Medicaid HMOs,” *Journal of Political Economy*, 2004, *112* (1), 1–47.
- Finkelstein, Amy, Matthew Gentzkow, and Heidi Williams**, “Sources of Geographic Variation in Health Care: Evidence from Patient Migration,” *Quarterly Journal of Economics*, 2016, *131* (4), 1681–1726.
- Gaynor, Martin, Kate Ho, and Robert J Town**, “The Industrial Organization of Health-Care Markets,” *Journal of Economic Literature*, 2015, *53* (2), 235–284.
- Grabowski, David C**, “The Economic and Clinical Outcomes of Nursing Home Care,” *National Bureau of Economic Research Working Paper*, 2006, (11893).
- MACPAC**, “Annual Analysis of Medicaid Disproportionate Share Hospital Allotments and Institutions for Mental Disease Expenditures,” Technical Report, Medicaid and CHIP Payment and Access Commission, Washington, DC 2024.
- Skinner, Jonathan**, “Causes and Consequences of Regional Variations in Health Care,” *Handbook of Health Economics*, 2011, *2*, 45–93.
- Stone, Robyn I and Joshua M Wiener**, “Who Will Care for Us? Addressing the Long-Term Care Workforce Crisis,” *Health Affairs*, 2001, *20* (4), 164–178.

U.S. Department of Justice and Federal Trade Commission, “Horizontal Merger Guidelines,” <https://www.justice.gov/atr/horizontal-merger-guidelines-08192010> 2010.

Wennberg, John E, *Tracking Medicine: A Researcher's Quest to Understand Health Care*, Oxford University Press, 2010.

A. Appendix

A.1 Stability of Geographic Patterns

To assess whether the geographic concentration patterns we document are stable over time or driven by particular subperiods, we compare ZIP-level spending in the pre-COVID period (January 2018 – February 2020) with the post-COVID period (January 2022 – November 2024). The Spearman rank correlation between ZIP-level spending across these two periods is 0.957, indicating that the geographic distribution of Medicaid spending is remarkably stable despite the disruptions of COVID-19, the ARPA HCBS spending increase, and the Medicaid unwinding. The ZIP codes that were the largest Medicaid billing centers in 2018 remain the largest in 2024.

This stability is consistent with the organizational structure hypothesis: because spending geography reflects the locations of established billing organizations rather than transient demand patterns, it changes slowly as organizational presence evolves. New fiscal intermediaries do not appear and grow to scale quickly; existing large organizations maintain their market positions over time.

A.2 Individual vs. Organization Providers

A key feature of the T-MSIS data is the distinction between individual providers (Entity Type 1 in NPES, representing sole practitioners and individual clinicians) and organizational providers (Entity Type 2, representing agencies, clinics, and intermediaries). In New York, organizations constitute 30.6% of billing providers but account for 96.7% of total Medicaid spending. This imbalance is more extreme than the national average and reflects the dominance of home care agencies and fiscal intermediaries in New York’s Medicaid landscape.

Individual providers, who represent 69.4% of billing NPIs, contribute only 3.3% of spending. These individuals include solo-practice therapists, nurse practitioners, and physicians whose Medicaid participation supplements their primary income. Their median tenure is shorter (18 months vs. 48 months for organizations), and their mean monthly spending is orders of magnitude lower. The coexistence of a large population of low-volume individual providers with a small population of high-volume organizations is consistent with a “hub-and-spoke” model where organizational intermediaries anchor the Medicaid billing system.

A.3 Data Coverage and Suppression

The T-MSIS dataset applies cell suppression when the number of unique beneficiaries in a provider–procedure–month cell falls below 12. This suppression removes approximately 63

million claims from the annual totals visible in the data, though the spending share removed is negligible because suppressed cells by definition involve few beneficiaries.

For New York specifically, the suppression rate is comparable to the national average. The practical consequence is that our ZIP-level analysis is most reliable for high-volume ZIP codes (where suppression is rare) and should be interpreted with caution for low-volume rural ZCTAs where individual provider–procedure cells may be suppressed. Our main findings—the dominance of T1019, the geographic concentration in Brooklyn, the high HHI values—are driven by the highest-volume cells and are therefore robust to suppression.

A.4 Robustness to Non-Spending Measures

Because the T-MSIS “Medicaid Amount Paid” field commingles fee-for-service payments with managed care encounter valuations, we verify that our core findings are robust to alternative quantity measures. The dominance of T1019 persists when measured by claim counts: T1019 accounts for 29.8% of all New York Medicaid claims, still the single largest procedure code. The geographic concentration is similarly robust: the Spearman rank correlation between ZIP-level total spending and ZIP-level total claims is 0.98, indicating that spending-based and claims-based geographic rankings are nearly identical. The organizational dominance finding (Type 2 NPIs account for 96.7% of spending) is even more extreme in claims: organizations generate 97.3% of all claims. These results confirm that the patterns we document reflect genuine concentration in service volume, not artifacts of managed care encounter valuation.

A.5 Specialty Composition by Region

Provider specialty composition varies substantially across regions, further illustrating the heterogeneity within New York State. NYC’s spending is dominated by home health agencies (57.3% of spending) and “Other” providers (20.4%, which includes fiscal intermediaries). By contrast, Upstate spending is more evenly distributed across assisted living (33.0%), clinic/center (10.3%), and physician (9.0%) categories. Long Island’s profile resembles NYC’s, with home health agencies capturing 47.3% of spending. Hudson Valley shows an intermediate pattern, with 34.2% home health and 16.1% assisted living.

These regional differences in specialty mix suggest that the within-state variation in Medicaid spending is not merely a matter of scale but reflects fundamentally different care delivery models. The NYC model is organized around home-based personal care delivered through organizational intermediaries; the Upstate model is organized around facility-based services (assisted living, nursing homes) and clinical providers.

Acknowledgements

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

Contributors: @SocialCatalystLab

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

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