The Evolution of Transit Access to Jobs in American Cities

Sierra Arnold*

Draft: September 16, 2026

Click here for the most recent version

Abstract

Households without a car are largely limited to the jobs reachable by public transportation. Whether that reach improved over the past decade, and whether any improvement came from transit service or from where jobs located, has not been measurable, because no consistent sub-annual series of job access by transit has existed for U.S. cities. We build a neighborhood-level, quarterly panel for 52 metropolitan areas over 2014–2025, routed from archived schedules. Job access by transit rose, but not because transit service improved. How many jobs a neighborhood can reach by transit depends on two things, the schedule and where the jobs are, and we separate the two. Tracking the same neighborhoods throughout, the typical neighborhood's job access by transit rose about 15% over the decade, collapsing below its 2014 level in 2020 before ending above its pre-pandemic peak. Decomposing that rise on the subset of metros whose archives are clean throughout, job growth alone would have produced more than the whole of it. The schedule's own contribution was roughly nothing before 2020 and negative after, a level we bound against agency-reported service data; job relocation

*Xavier University. Contact: arnolds7@xavier.edu. Archived transit schedules were obtained from Transitland (`https://www.transit.land/terms`), whose Feed Archive and academic download allocations made the panel possible, and from the Mobility Database maintained by MobilityData; Philadelphia's 2021–2025 schedules were obtained from SEPTA's published GTFS archive. The schedules themselves are produced by the transit agencies of the metropolitan areas we study, whose open-data publication this research depends on. Street network data are © OpenStreetMap contributors, available under the Open Database License (`https://www.openstreetmap.org/copyright`), via Geofabrik extracts. Employment data are from the U.S. Census Bureau's LEHD Origin–Destination Employment Statistics; demographic context is from the American Community Survey. The 2010–2020 block-group crosswalk is from IPUMS NHGIS (Schroeder et al., 2026). Agency-reported service data are from the Federal Transit Administration's National Transit Database. Published *Access Across America* values used in the validation are from the University of Minnesota's Accessibility Observatory. Routing was performed with r5py (Fink et al., 2022) over Conveyal's R5 engine (Conway et al., 2018), on resources provided by the Ohio Supercomputer Center (Ohio Supercomputer Center, 1987), whose support we gratefully acknowledge. Any errors in the assembly or interpretation of these data are our own.

toward transit-served places returned most of what the schedule took away. The gains also ran opposite the level gradient. Access levels favor poor, central, car-free neighborhoods, but the growth landed in richer, car-owning, peripheral neighborhoods where the jobs sprawled. The poorest quartile gained essentially nothing, and majority-Black neighborhoods lost longer-reach access outright once archive entries are removed (−13% within 60 minutes), in the handful of metros where their population concentrates. Weighted by car-free households, the rise shrinks to +2.5%, against +15% for the average resident. Within cities, at the same distance from the same core, growth accrued to richer neighborhoods, overwhelmingly through employment location. The paper is descriptive; we make no causal claims about employment.

Keywords: Job accessibility; public transit; GTFS; spatial mismatch; transport equity; accessibility decomposition

1 Introduction

Transit determines which jobs a household can reach without a car. In a typical U.S. metropolitan area a worker can reach far more jobs by driving than by riding transit within the same travel-time budget, a modal access penalty and spatial mismatch (Kain, 1968; Gobillon et al., 2007). The penalty binds hardest where driving is not an option: about 8% of U.S. households have no vehicle available, and 19% of lower-income and 21% of Black adults seldom or never drive, compared to 3% of upper-income and 7% of White adults (Pew Research Center, 2024; U.S. Census Bureau, 2023). For low-income neighborhoods with few cars, access by transit is the opportunity set. This paper asks how that opportunity set changed across American neighborhoods between 2014 and 2025. Two questions follow from the first: whether the change was the work of transit service or of the shifting location of jobs, and which neighborhoods it reached.

The decade from 2014 to 2025 spans the COVID-19 pandemic, the largest shock to American transit in generations: ridership collapsed, agencies cut service, employment patterns shifted, and downtowns emptied. Little is known about how these changes altered access to employment. What a neighborhood can reach depends on the timetable and on where employment locates, and the decade moved both. Whether Americans ended it better or worse connected to jobs by transit cannot be answered from ridership counts or service budgets alone.

That job access by transit is low, and the transit/car gap large, is well established. An annual national series reports levels (Owen et al., 2025), a block-group program maps the transit/car gap (Maharjan et al., 2024), a pandemic-era dashboard tracks seven regions from 2020 onward (Klumpenhouwer et al., 2021), and a recent single-city panel follows New York across four snapshots (Zhang et al., 2026). None yields a national, sub-annual, decade-spanning panel of neighborhood job access by transit; without one, a measured change cannot be told apart from a change in method or data vintage, and the network's and the jobs' contributions to a real change cannot be separated.

We build that panel from archived GTFS schedules, routed on one fixed method four times a year for 52 metros over 2014–2025, with New York, an order of magnitude more transit-rich and with gradients that invert the national ones, analyzed separately (Section 6). The typical neighborhood's job access by transit rose about 15% over the decade, collapsing below its 2014 level in 2020 before recovering above its pre-pandemic peak. The improvement was largely produced by labor-market geography rather than transportation investment: metro job growth alone would have delivered the entire rise, the network's own retreat took part of it back, and the relocation of jobs toward transit-served places returned most of what the network subtracted. And the decade's gains ran opposite the level gradient, landing in richer, car-owning, peripheral neighborhoods: weighted by car-free households, the rise is +2.5%, against +15% for the average resident. The transit/car gap, an order of magnitude or more, we quantify as an upper bound and read cross-sectionally rather than track over time, since its movement is dominated by the car denominator rather than by transit.

Section 2 reviews the accessibility and spatial-mismatch literatures the paper builds on. Section 3 describes the schedule archive, how the panel is built from it, and how we audit it for archival artifacts. Section 4 defines the access measure, the decomposition, and the inference. Section 5 reports the results: what drove the rise, the trajectory in full, who has access and who gained it, the transit-reliant margin, the metros, and the pandemic. Section 6 treats New York on its own. Section 7 discusses what the panel does and does not establish. Appendix A holds the supporting exhibits.

2 Related Literature

Accessibility measures the opportunities that can be reached from a location within a given travel budget, a concept dating to Hansen (1959). Subsequent work has developed cumulative-opportunity, gravity-based, and utility-based measures and has clarified the strengths and limitations of each approach (Handy and Niemeier, 1997; Geurs and van Wee, 2004). In practice, however, cumulative-opportunity and gravity-based measures generally produce similar conclusions when calibrated to comparable travel-time thresholds (Santana Palacios and El-Geneidy, 2022), making the availability of consistent longitudinal data more consequential than the particular accessibility metric employed. Accessibility fluctuates within the day, though a constant peak-hour measure tracks relative accessibility across neighborhoods closely (Boisjoly and El-Geneidy, 2016). A further methodological thread decomposes a change in accessibility into the part attributable to the transport network and the part attributable to the population or land-use (here, job) distribution (Stepniak and Rosik, 2018; Rosik et al., 2025). A related literature separates the same two contributions across design scenarios (Lahoorpoor et al., 2022). Applying these decompositions to observed changes over time, however, requires accessibility measures computed consistently across years, which is the requirement the panel developed here is built to meet (Section 3).

While this literature establishes how accessibility is measured, the more important question for this paper is what those measures reveal about U.S. transit systems. Applied to U.S. transit, these measures have produced a body of direct evidence on how little of a metropolitan job market a rider can reach. Tomer et al. (2011) provide a 100-metro cross-section, and the *Access Across America* series (Owen et al., 2016), building on continuous-accessibility methods (Owen and Levinson, 2015), reports transit accessibility annually, with a parallel automobile series. Its recent editions report year-over-year comparisons (Owen and Murphy, 2020) and describe a decade of consistently calculated data (Owen et al., 2025). Because each edition pairs that year's schedules with that year's employment data, its year-over-year changes combine service change with economic change. The Transit Equity Dashboard (Klumpenhouwer et al., 2021) tracks block-group transit access and its demographic incidence in seven large U.S. regions from February 2020 forward. Closest to our object is the *Modal Access Gap* program: transit versus car job access computed at the block-group level across the largest U.S. metropolitan areas, with an explicit focus on carless and underserved populations (Maharjan et al., 2024). The same group measures transit against walking in the 50 largest metros, reading the gap as an effectiveness indicator (Ermagun and Witlox, 2024). This work establishes the cross-sectional gap and its spatial structure, including the finding that carless households tend to live where the gap is narrower; its temporal dimension is travel time rather than calendar time (Janatabadi et al., 2023). Consequently, existing studies can describe where transit accessibility is high or low, but they cannot distinguish how observed changes over time reflect changes in transit networks versus changes in the spatial distribution of employment.

What gives these magnitudes their economic weight is the spatial mismatch hypothesis (Kain, 1968), which holds that physical disconnection from employment depresses labor-market outcomes for disadvantaged workers. The hypothesis and its mechanisms have been reviewed extensively (Ihlanfeldt and Sjoquist, 1998; Gobillon et al., 2007). A strand of this work stresses that mismatch is as much modal as spatial: a worker without a car faces a far smaller opportunity set than a driver from the same location (Grengs, 2010). Kain's hypothesis also has a job side. Employment in U.S. metros has decentralized for decades

(Glaeser and Kahn, 2001), with job share shifting away from the urban core in nearly every large metro through the 2000s (Kneebone, 2009). That movement is the one Kain originally identified as the mismatch's engine. Whether decentralization has been pulling jobs toward or away from the places transit serves is an empirical question about the joint evolution of the two, and answering it requires observing both over calendar time.

How access is distributed across people and places has been studied on both empirical and normative grounds; the normative case is reviewed by Pereira et al. (2017). Foth et al. (2013) track transit access against neighborhood social need in Toronto across two census years, and Yeganeh et al. (2018) compare access by income and race across U.S. metropolitan areas in a single cross-section. Most closely, Zhang et al. (2026) track block-group transit access to jobs and services across four snapshots for New York City, decomposing a widening racial access gap into transit-network and residential-sorting components. The measurement is made feasible by general-purpose routing over open schedule data (GTFS) and street networks (Conway et al., 2018; Pereira et al., 2021). The closest precedent reconstructs a retrospective timetable from vehicle-location data (Wessel et al., 2017). The pandemic's effect on transit service and ridership is documented at the system and metro level (Liu et al., 2020; Kar et al., 2022), an episode within the decade our panel covers rather than its subject.

Taken together, these literatures establish the level of U.S. job access by transit, the cross-sectional shape of its modal gap, and its distribution at a point in time, and they supply the machinery for decomposing a change in access. The object that remains open is job access by transit observed in calendar time: at the block group, on a single method held fixed, at sub-annual resolution, across many metros. That is the panel this paper builds, and it draws on each of these literatures in turn. We adopt the transparent cumulative measure, carrying gravity and logsum variants as robustness, and allocate every change between the network and the jobs with a reference-period-invariant Shapley decomposition (Shorrocks, 2013). *Access Across America* serves as a validation benchmark and the National Transit Database as an independent check on the network component (Appendix A). The modal gap enters as a secondary, cross-sectional upper bound, and the strict-contemporaneity and service-calendar rules of Section 3.2 meet the archived-feed fidelity issues directly. Throughout we measure access, not employment; for evidence on whether transport access affects labor-market outcomes see Sanchez (1999), Tyndall (2017), and the meta-analysis of Bastiaanssen et al. (2020).

3 Data

We assemble a neighborhood-level panel of transit and car job access for 52 U.S. metropolitan areas outside New York, plus New York itself, measured four times a year for the twelve years from 2014 through 2025, from four ingredients: an archive of transit schedules, the location of jobs, the street network, and demographic context. Assembling consistently routable schedules at sub-annual resolution, on harmonized dates across more than a decade, has not been done before.

3.1 Data sources

Transit service is described by General Transit Feed Specification (GTFS) feeds. GTFS defines a common format for public transportation schedules and the geographic information that accompanies them: an agency publishes its routes, its trips, the location of every stop, and the time a vehicle is scheduled to serve each one. The specification exists so that this data can be consumed by software, which is what puts transit directions into applications like Google Maps, and nearly every transit agency in the country adheres to it. There is, however, no national repository of feeds. Schedules live with the agencies that publish them, and superseded versions are ordinarily discarded, so a feed downloaded today describes only today's service. What makes a historical panel possible is that community archives capture dated snapshots and keep them. We assemble ours from two such archives, Transitland and the Mobility Database, supplemented by SEPTA's own published archive for Philadelphia's recent years.

Job locations come from the Census LEHD Origin–Destination Employment Statistics (LODES) Workplace Area Characteristics, which report job counts at the workplace Census block annually; we aggregate them to the block group, so that both ends of a trip are measured on the same geography: we route from block-group origins and count jobs at block-group destinations. We count all jobs throughout, and version 8 of LODES reports every year on 2020 census blocks, so the destination geography is stable across the panel. Because the jobs data are annual, sub-annual movement within any year is driven entirely by changes in the transit schedules.

Door-to-door routing uses the OpenStreetMap (OSM) street and path network. This maps the walk-ride-walk time for transit, and the drive time for cars. The car alternative is free-flow using posted speeds with no congestion, parking, or access-and-egress penalty, which overstates real driving access and hence the size of the gap. The gap's level is thus best read as an upper bound on the car advantage. A present-day OSM extract is used, so the car baseline reflects today's roads. The quarterly trajectory is driven by the transit

Table 1: Descriptive statistics, 52-metro balanced panel outside New York.

	Mean	SD	p10	p50	p90
A. Job access (block group × quarter)					
Transit jobs ≤ 45 min	44,214	115,756	276	6,768	107,221
Car jobs ≤ 45 min	1,536,484	1,204,558	237,127	1,181,697	3,163,993
Transit/car ratio R	0.034	0.120	0.0005	0.0058	0.100
B. Neighborhood characteristics (block group)					
Population	1,500	800	686	1,374	2,455
Median household income ($000)	96.7	49.1	43.5	86.8	164.4
Below 200% poverty (%)	27.0	20.4	4.8	22.0	57.2
Zero-vehicle households (%)	7.7	11.7	0.0	3.2	21.9
Non-Hispanic White (%)	52.1	30.7	6.0	56.5	90.8
Non-Hispanic Black (%)	12.9	21.8	0.0	3.2	42.1
Hispanic (%)	23.0	25.5	0.2	12.9	65.1
Non-Hispanic Asian (%)	7.3	12.1	0.0	2.3	21.4
Distance to urban core (km)	7.2	11.0	0.6	3.4	18.3

Panel A is pooled over block-group × metro × quarter (flagged cells excluded), 45-minute budget, all jobs; R is the transit/car ratio, an upper bound (Section 4). Panel B is per block group ($n \approx$ 85,700), from ACS 2018–2022. New York excluded.

series.

Neighborhood income, race, and vehicle availability come from the American Community Survey, used both to characterize who has car access and to identify transit-reliant neighborhoods (high zero-vehicle households and high below-poverty shares). This is a neighborhood-level intersection rather than a household-level one. The block group is the unit throughout; 2020 block-group population centroids anchor each neighborhood's origin for routing. Table 1 reports descriptive statistics for the access measures and the neighborhood characteristics.

3.2 Panel construction

For each operator serving a metro we take the feed in effect on the second Wednesday of January, April, July, and October of each year (a representative mid-week, mid-quarter service day), combine all the operators that serve the metro into a single routable network, and route that schedule. A metro-quarter is thus the service that was actually operating on that date. Seasonal timing does not drive annual comparisons: quarter-of-year effects are small, each quarter sitting within about a log-point of its year's mean, so nothing of substance hinges on the second-Wednesday date.

Agencies republish their feeds at every service change, several times a year in the recent era, and these versions are archived, so sub-annual sampling resolves service changes, above all the COVID-19 collapse and recovery, that an annual snapshot averages away. This sub-annual depth is feasible only in the recent era; routed quarters exist back to 2013, but the archive there is thin, so the analysis window, and every exhibit, runs 2014–2025. Second, no substitution. A metro-quarter includes an operator only when its feed is strictly contemporaneous, meaning its service period spans the sample date; an operator without one is dropped from that quarter rather than backfilled with a stale feed. Because we drop rather than impute, coverage gaps shrink the sample but do not bias the cells that remain.

Coverage in the archive is uneven (sparse in the early years, densest in 2016–2019, and thinner through a 2020–2023 trough), so the analysis uses a near-balanced panel: the 52 metropolitan areas outside New York (Fig. 1) present in at least 36 of the 48 quarters of 2014–2025. The composition-adjusted trend is then not contaminated by metros entering and leaving the sample. We refer to it as the balanced panel for brevity; per-metro coverage is in Appendix A. New York is set aside as a separate case, because its extraordinary transit supply both dominates any population-weighted average and inverts gradients found elsewhere. This panel is the sample every level, trajectory, and distributional exhibit uses, narrowed to the metros where both endpoint windows are observed (45 of 52) only when an exhibit needs a decade change; the decomposition alone uses a stricter cut, the clean-archive core, defined where it is used (Section 5.1). Table A.2 lists each exhibit's sample.

Figure 1: The 52 panel metropolitan areas outside New York

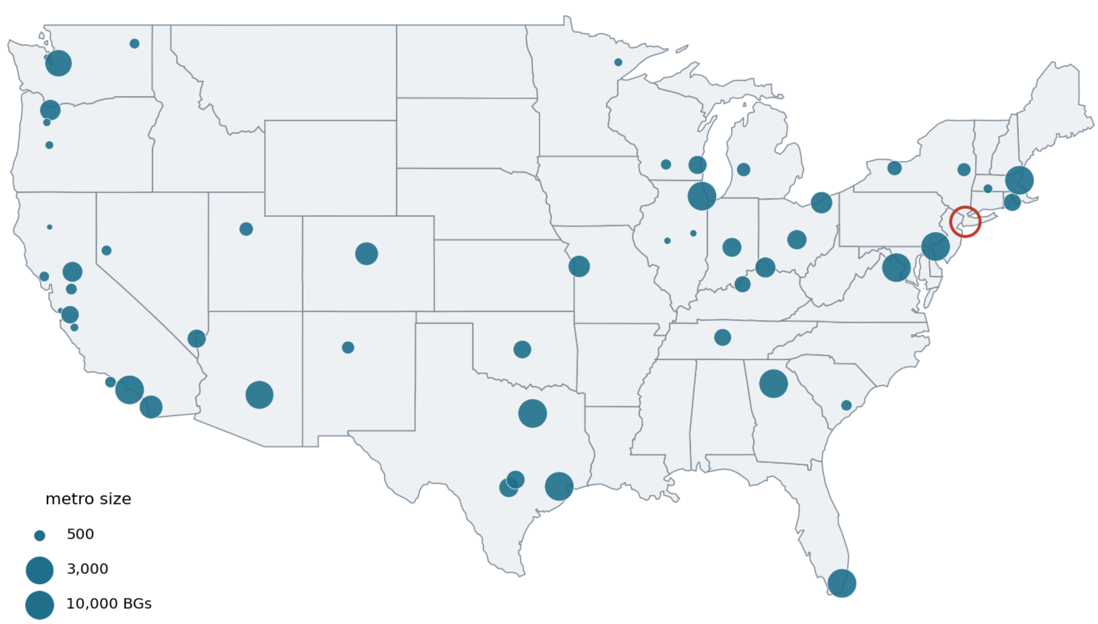

The 52 U.S. metropolitan areas outside New York analyzed at quarterly resolution, 2014–2025; point size is proportional to the number of block groups. New York, analyzed separately in Section 6, is the hollow marker.

3.3 Data quality and representativeness

An operator whose feeds were never archived in the early years is invisible to quality scans, which detect changes rather than absences; this archive adoption is removed by the trajectory's archive-entry correction (Section 4) rather than left to contaminate the trend. Feeds that are archived can still misrepresent the service actually running; we flag the affected metro-quarters rather than dropping them, and flagged cells are excluded from all value figures while still counting as present for coverage (Appendix Table A.3 lists the classes, detection rules, and counts). Because internal scans cannot catch everything, the archive is also corroborated against agency-reported service, which flagged cells every internal check had passed (Appendix A). Finally, GTFS encodes scheduled rather than delivered service (Section 7).

The panel over-represents larger, transit-intensive metropolitan areas, because inclusion depends on historical GTFS availability rather than random sampling: the 52 cover about 45% of total U.S. metropolitan population, including 30 of the 50 largest metros and 40 of the largest hundred, and are more transit-oriented, though no more car-free, than the typical metro, while missing several of the largest legacy systems whose archives do not span the window (San Francisco–Oakland, Minneapolis–St. Paul, Detroit, Baltimore; Washington is included, but its WMATA feed is unarchived, so its access is understated throughout). The omissions are plausibly one-signed: the legacy systems that are in the panel are its decliners (Section 5.6), so the absent ones, large, old, and retrenching after 2020, would more likely pull the trajectory down than up. The full panel-versus-nation profile, per-metro coverage, and the named omissions are in Appendix A (Tables A.1, A.4). We read the results as describing job access by transit in the covered metros rather than in the average U.S. metro.

4 Measuring Transit Access to Jobs

4.1 The accessibility measure

The unit of the panel is the block group × metro × quarter. For each metro-quarter we build a routable network from the metro's OpenStreetMap extract and the operators' strictly-contemporaneous GTFS feeds and compute door-to-door travel times between every pair of block-group population centroids.

Travel times are computed with R5 (Conway et al., 2018). For transit we depart across a two-hour morning-peak window (7:00–9:00) and take the median door-to-door time over departures within it (the initial walk to a stop, in-vehicle time, waiting and transfers, and the final walk to the destination), so that a value reflects typical rather than best-case timing and is not hostage to a single fortunate departure minute. Trips longer than 90 minutes are treated as out of reach; the 45-minute primary budget sits far from that censoring point, so neither the ceiling nor the median-over-window convention is load-bearing. Walking legs use the engine's default walking speed with no separate distance cap (walking is bounded by the time budget itself), and all remaining routing parameters are the engine defaults, held fixed across every metro and quarter. Because routing runs centroid to centroid, door-to-door times are coarsened in large peripheral block groups; the centroids are fixed, so this measurement error is time-invariant and nets out of the within-neighborhood trajectory. The car alternative is produced by the same engine on the same street network, substituting driving for the transit-and-walk legs, so the two modes differ only in how a traveler moves and not in where origins, destinations, or roads sit. From the resulting travel-time matrix we form cumulative access: the count of jobs at destinations reachable within a fixed time budget. The cumulative job access by transit of origin block group i in metro-quarter q at budget τ is

$$A_{i,q}^{(\tau)} \;=\; \sum_j \mathrm{jobs}_{j,y(q)} \cdot \mathbf{1}\left[\, t_{ij}^{\mathrm{transit}}(q) \le \tau \,\right], \qquad \tau \in \{30, 45, 60\} \text{ min},$$

where $t_{ij}^{\mathrm{transit}}(q)$ is the transit travel time on the schedule operating in quarter q and $\mathrm{jobs}_{j,y(q)}$ is jobs at destination j in the annual LODES vintage $y(q)$ of that quarter; car access $A_{i,q}^{(\tau),\mathrm{car}}$ is identical but for car travel times on the fixed street network.

We report the 45-minute budget as primary and the full 15/30/45/60-minute curve alongside it; we count all jobs (LODES C000) throughout (Section 3.1 explains why no low-wage band is reported). Cumulative access is the most transparent and most widely reported construct; gravity and utility-based (logsum) measures, which discount farther destinations, are about 0.9 correlated with it once calibrated to a common budget (Santana Palacios and El-Geneidy, 2022), so we carry them as robustness (Table A.6, Appendix A). Access climbs steeply with the time budget, roughly forty-five-fold between 15 and 60 minutes, but the budget series move in parallel across the decade, so nothing in what follows hinges on the 45-minute choice; the full 15/30/45/60-minute curve is Appendix A.

The primary object is job access by transit itself; Fig. 2 shows what it looks like on the ground for one mid-size panel metro: the characteristic access mountain, peaking downtown and falling roughly two orders of magnitude to the periphery. We summarize it population-weighted (the access of the average person, the convention in the cumulative-opportunity literature, Owen and Levinson 2015) and report the median block group, the typical neighborhood, as a companion. Access is sharply right-skewed within a metro, near zero across the low-service periphery and high in the service-bearing core; in the metro mapped here the median block group reaches about 5,400 jobs.

Figure 2: Job access by transit per block group: Cincinnati, 2023Q4

Jobs reachable within 45 minutes by transit, per block group, Cincinnati 2023Q4 (log scale, central 30 km). Every panel cell is one such map, per metro per quarter.

4.2 The transit/car comparison

The transit/car gap enters only as a secondary comparison: the ratio $R_{i,q} = A^{(\tau),\mathrm{transit}}_{i,q} / A^{(\tau),\mathrm{car}}_{i,q}$ at the same origin (so $R = 0.01$ is a hundredfold car advantage). We read the gap as an upper bound, since the free-flow car baseline overstates driving access, and we do not track it over time: its movement is dominated by the car denominator and by sprawl rather than by transit. The transit-access trajectory, not the gap, carries the temporal narrative; the gap's magnitudes are reported with the levels results (Section 5.3) and by time budget in Appendix A.

4.3 Measuring change

We read how access moved two ways. The first is a composition-adjusted trajectory: log access regressed on quarter indicators with a block-group $\times$ archive-regime fixed effect. The regime term is the archive-entry correction, the splice familiar from chained price indices: because the archive itself expanded (every panel metro adds at least one operator to its archived set after its first year), an operator's entry into the archive would otherwise read as new service. Growth is therefore measured only while a metro's archived operator set is unchanged, with the level step at each change absorbed rather than counted (Fig. 3). The correction drops no observations, and it is conservative: genuinely new agencies are removed too, and treating the five operator exits, the one anti-conservative case, as real changes moves the 2025 endpoint by less than a tenth of a log-point. We adjudicated the large regime-boundary steps against agency-reported service: the adjudicable cases were archive events rather than service change, the exception being Columbus, whose entry is real new service.

Figure 3: The archive-entry correction in one metro: Miami

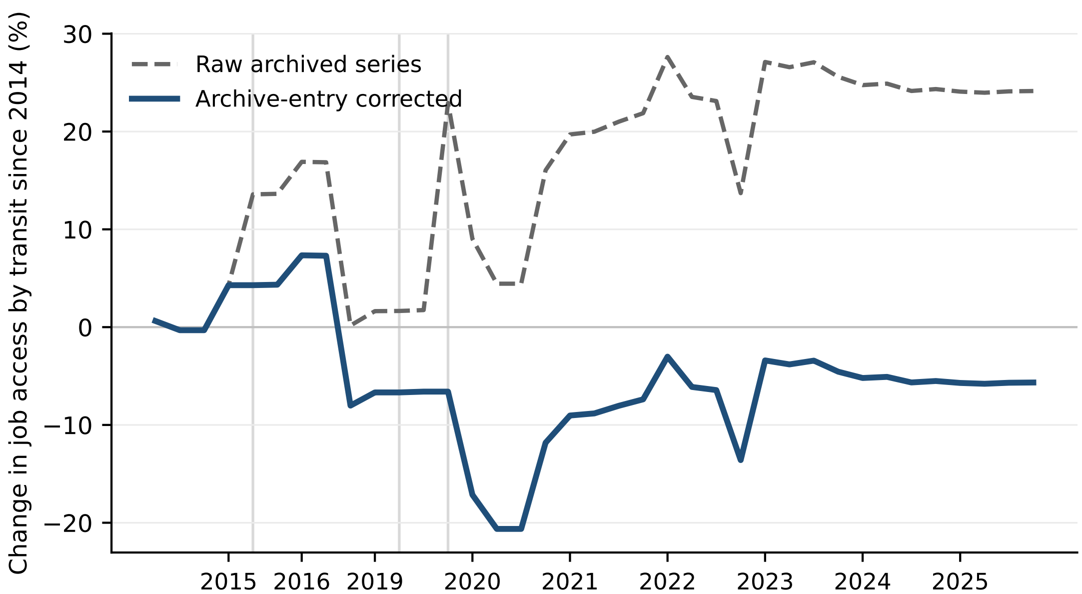

Population-weighted mean job access by transit in Miami: the raw archived series (dashed) and the series with the archive-entry correction applied per block group (solid), percent change against the 2014 annual mean, over the 3,942 block groups present in every covered quarter (2017–2018 are uncovered; the axis spaces covered quarters evenly). Vertical lines mark changes in Miami's archived operator set; the separations between the two series at those marks are the archive-entry steps the correction removes.

Every value in the panel, each metro, quarter, and mode, is produced by a single procedure held fixed, so a difference in access between two quarters reflects a difference in service (or, across years, in the location of jobs), not a change in measurement. This is what lets us read the COVID-19 trajectory (Section 5.7) as a real movement rather than an artifact of how access was counted.

4.4 The decomposition

The second reading is a network-versus-jobs decomposition of each neighborhood's log change. Conceptually, we ask what access would have looked like if only transit had changed, and separately if only jobs had changed. Holding the 2014 job map fixed and letting only the network change isolates the network's own contribution; holding the 2014 network fixed and letting only jobs move isolates the job side. In logs the two contributions sum exactly to the observed change at every neighborhood; we allocate them symmetrically, following the transport-versus-land-use decomposition thread (Rosik et al., 2025) with the symmetric allocation of the Shapley construction (Shorrocks, 2013). The job side splits exactly in two: metro-wide job growth, which involves no geography, since cumulative access is linear in the jobs vector and uniform growth enters as the growth of the metro's job total; and job relocation, employment shifting toward or away from the places transit serves. The decomposition moves on annual steps, the cadence at which both components can move, because LODES jobs are an annual vintage frozen within each year; components are indexed to the 2014 annual average, since the archive's early base quarters are thin (Section 5.1).

4.5 Inference and validation

For inference we cluster standard errors at the block group, which absorbs the serial correlation of a neighborhood with itself over time. By the design-based logic of Abadie et al. (2023) this is the right uncertainty for describing the panel we observe, since we hold essentially the entire population of its block groups. For claims that read as a common process across metros, we additionally report metro-clustered intervals (Table A.7, Appendix A) and scope the text to whichever the evidence supports. Metro job growth, job relocation, and the net decade change are metro-robust; the transit-network component and the racial growth contrasts are statements about this panel, whose metros genuinely diverge.

Two external checks anchor the measure. Across the 390 metro-years overlapping the published *Access Across America* series, our levels agree at log correlations of 0.88–0.95 (Appendix A). Our levels sit at a stable fraction of theirs, as expected: they weight by workers at their residence over the urban area, we weight by population over the larger cropped metropolitan area, whose low-access periphery pulls our mean down, and departure-time protocols and employment vintages also differ. The informative comparison is therefore cross-metro agreement rather than equality of levels. And the measure predicts behavior: metros with more transit-reachable jobs are metros where more workers actually commute by transit (Fig. 4).

Figure 4: Metro job access by transit against the ACS transit commute share

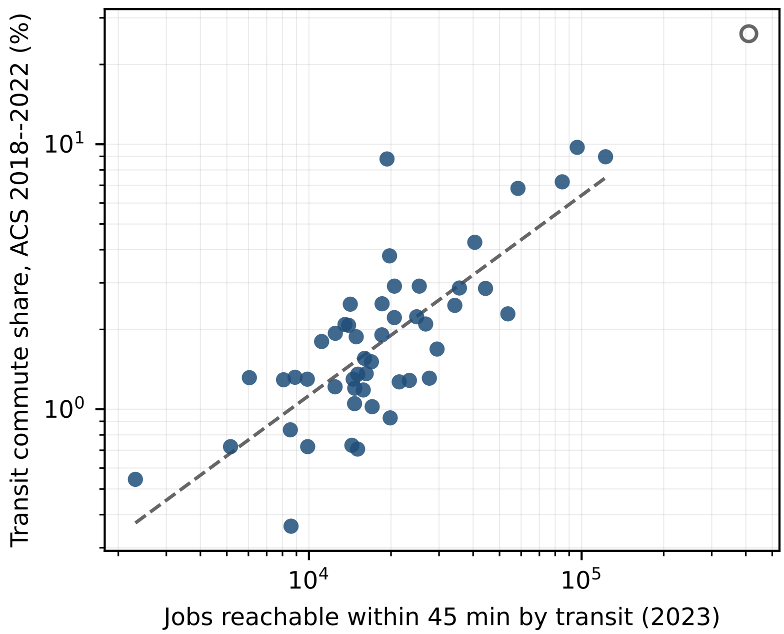

Metro job access by transit (2023, flagged cells excluded) against the ACS 2018–2022 transit commute share (table B08301), 51 metros, log scales; the dashed line is the ex–New-York fit (slope ≈ 0.8; $r = 0.77$, or 0.82 including New York, shown hollow).

5 The Evolution of Job Access by Transit

Unless noted, results are on the panel, the balanced 52 metros outside New York (Section 3.2), population-weighted, at the 45-minute budget for all jobs; the decomposition alone uses the clean-archive core defined there. Trajectories are composition-adjusted by within-block-group fixed effects, and confidence bands are clustered at the block group (Section 4). New York, an order of magnitude more transit-rich and with several gradients that run the other way, is analyzed on its own in Section 6. The section moves from the national arc, what drove the rise and the trajectory in full, to the distribution: within cities, where the panel's neighborhood resolution is decisive, we ask who has access, who gained it, and how the transit-reliant margin the introduction motivates fared; then across metros, closing with the pandemic episode.

5.1 What drove the rise: network or jobs?

Over 2014–2025 the typical neighborhood's job access by transit rose about 15%. But the more revealing question is what drove that rise: did job access by transit improve because the network expanded, or because the jobs moved toward it? The decomposition of Section 4 separates the two.

The answer is unambiguous and is the paper's central finding (Fig. 5). The decomposition is estimated on the clean-archive core: of the panel's 52 metros, the 35 whose block groups are observed unflagged in every year, and among those the 19 with no flagged quarter at all across the decade (28,147 block groups; membership in Table A.2). Metros with flagged eras are excluded whole, because their component paths carry archive quality drifting in and out rather than service; their measured networks "improve" through the pandemic. The exclusion protects the component paths: the 2025 endpoints are robust to it within a fraction of a log-point, and the jobs-side components are nearly identical on the full every-year sample (Fig. A.2). Measured against the 2014 annual average, and in the log-point units in which the components add, metro job growth alone would have added +15.9 log-points to the typical neighborhood's job access by transit as of 2025. The transit network's own contribution holds within a point of zero through 2019 and then falls sharply after the pandemic, ending at −6.4 log-points: the documented post-2020 service retrenchment, now visible as a component of access rather than as service hours. Job relocation contributes +5.7 log-points: employment did drift toward transit-served locations, recovering most of what the network's retreat took away. Where that relocation accrues resolves an apparent tension with the incidence results below: by distance to the core, the relocation component is slightly negative in the two inner quartiles (about −1 log-point) and strongly positive at the fringe (+16), concentrated in fast-growing metros. Employment moving outward exits the core's reachable set and enters the reachable sets of mid-ring and peripheral neighborhoods. The same outward drift that is sprawl from the core's vantage is, in the aggregate, relocation toward transit-served places. The three sum exactly to the observed +15.2 log-points, a rise of about 16.4%, on this sample. The 52-metro trajectory puts the decade at +14.1 log-points (+15.1%); Table A.8 reconciles the two, and the largest of its three steps is the sample itself: the clean-archive metros grew faster. That step moves the observed rise, not the split. The robustness stated above, components within a fraction of a log-point, holds the sample fixed and varies the flag treatment. Metro job growth and job relocation survive clustering at the metro level, and so does the net decade change (Appendix A): they are common processes across the 52 metros. The transit network component, by contrast, has a metro-clustered interval that spans zero; metros' network paths genuinely diverge (Section 5.6). Neighborhoods reach more jobs by transit than in 2014 because employment grew, not because transit improved.

Jobs are observed through 2023, so 2024–25 movement is network-only by construction; a decomposition ending at 2023 reads nearly identically ($+14.9 = +15.9 + 5.7 - 6.6$).

Figure 5: Network, job-growth, and job-relocation components of the change in job access by transit, 2014–2025

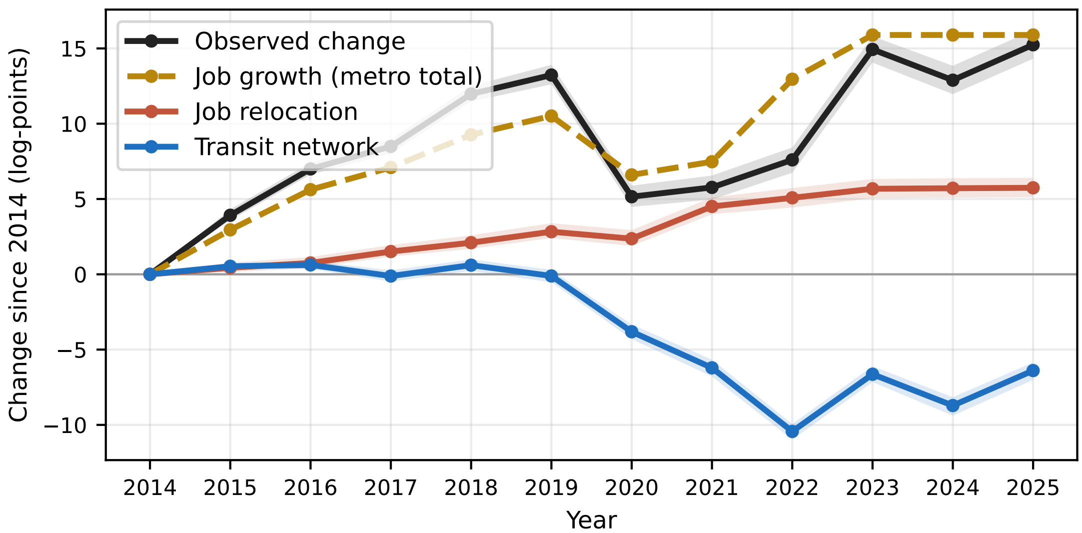

Decomposition of the change in job access by transit into the transit network (jobs held at 2014), metro-wide job growth, and job relocation components, within block group, annual, population-weighted, the clean-archive core (19 metros outside New York), indexed to the 2014 annual average; vertical axis in log-points, the units in which the three components sum to the observed change. Shading is the 95% block-group-bootstrap band.

5.2 The trajectory, in full

The decomposition summarizes an arc worth seeing at full resolution, and on the full 52-metro panel (Fig. 6). On the corrected index (Section 4) the typical neighborhood's job access by transit climbed about +9% by late 2019 (+8.9 log-points $[8.1, 9.7]$), collapsed during 2020 below its 2014 level (-3.7 $[-4.5, -2.9]$ in 2020Q3; the pandemic erased the entire decade's gain and then some). It recovered to about +15% by 2025 (+14.1 log-points $[13.3, 14.8]$; +15.1% in ratio terms), above its pre-pandemic peak. The uncorrected index ends far higher, at about +22 log-points (+25%); the roughly eight-point wedge measures how much of that raw climb is archive adoption rather than service. These intervals describe the panel; as a statement about a common national process the net decade change holds up, with a metro-clustered interval excluding zero (about $[+6, +24]$; Appendix A), though metros vary widely around it, the dispersion Section 5.6 documents. The sub-annual panel resolves the COVID-19 collapse and recovery at the neighborhood level, which an annual series cannot.

The raw population-weighted level, by contrast, is an unreliable guide to how access moved: it depends on how neighborhoods are weighted, on which metros are present in each period (the panel's composition shifts), and on the feed-quality artifacts the integrity scan flags (Section 3.2, Section 5.6). Across these choices it ranges from modestly positive to slightly negative. The within-block-group index compares each neighborhood to itself on a consistent set observed throughout, sidestepping all three, and is the trajectory of record. Including flagged cells, the corrected index ends at +17.7% rather than +15.1%; Appendix Table A.3 lists the flag classes, their detection rules, and the cells each caught.

Figure 6: Change in job access by transit relative to 2014, quarterly, 2014–2025

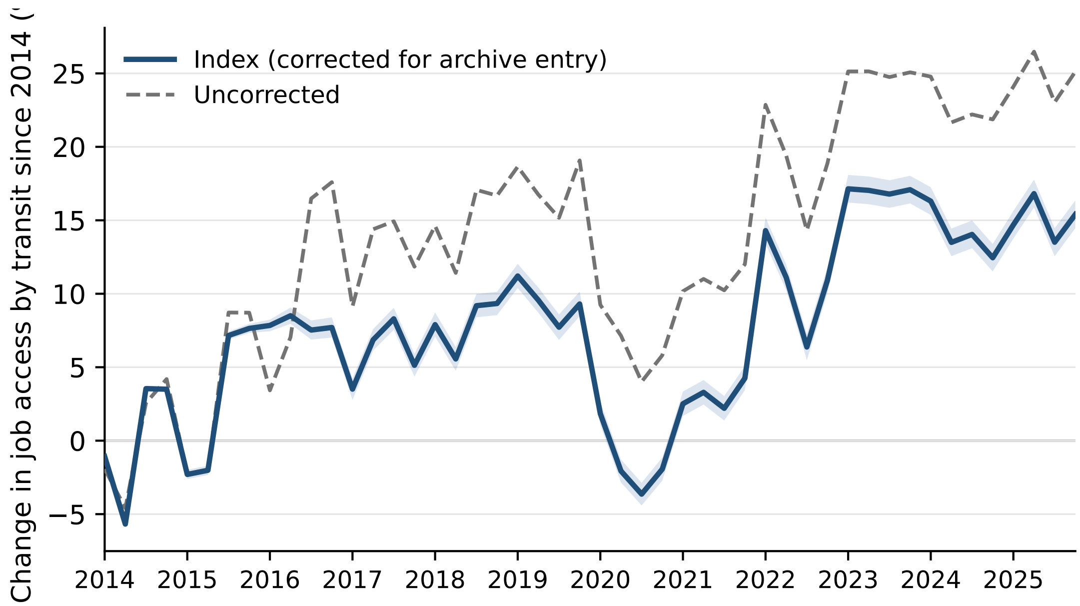

Within-block-group fixed-effects change in jobs reachable within 45 minutes by transit, relative to the 2014 average, population-weighted, balanced 52-metro panel outside New York. The solid line is the archive-entry-corrected index; the dashed line is the uncorrected series. Shading is the 95% block-group-clustered band.

5.3 Within cities: who has job access by transit?

The panel's neighborhood resolution makes the within-city distribution of access, and of its change, a first-class object, and it is where this panel can go that neither an area-aggregated series nor a single cross-section can.

The modal magnitude behind the comparison is large. Outside New York the population-weighted average neighborhood reaches roughly thirty-eight-fold fewer jobs by transit than by car within 45 minutes; for the median neighborhood the shortfall is about 172×; by time budget the gap runs 86× at 15 minutes and eases, without closing, to about twentyfold within an hour (Appendix A). The two summaries differ because access is severely right-skewed: the ratio of population-weighted means is pulled toward the transit-rich core, where transit fares comparatively well, while the median neighborhood's own ratio is far larger. When the abstract says the gap is "an order of magnitude or more," the order of magnitude is the smaller, mean-based figure.

Two questions with opposite answers, taken in turn. In levels (as of 2023), job access by transit is sharply graded toward the neighborhoods that rely on it. Population-weighted on the balanced panel, the poorest within-metro income quartile reaches about 1.9× the access of the richest; neighborhoods where more than a third of households are car-free reach about 14× the access of car-owning ones; high-poverty neighborhoods reach about 2.4× low-poverty ones; and central neighborhoods reach roughly 50× the urban fringe. Outside New York, majority-Black neighborhoods reach more than majority-White ones (about 2.1×), while majority-Hispanic neighborhoods reach about the same (1.0×). Each ratio's confidence interval excludes one, except the Hispanic contrast, whose interval spans it: genuine parity (Table 2). A natural objection is that these gradients merely restate urban form: poor, car-free, and minority households live centrally, where transit is. Controlling for distance to the urban core separates the two. The car-reliance gradient survives within every distance band: car-free neighborhoods reach more jobs by transit than their car-owning neighbors at the same distance from the center. The income and poverty gradients also survive within every band, strongest in the 5–30 km suburban ring and attenuated but still pro-poor in the dense core. The racial gradient, by contrast, is near one within distance bands: the majority-Black advantage in the pooled comparison reflects where those neighborhoods sit (closer to the core) rather than race per se. We report it as such.

Table 2: Job access by transit across neighborhood groups.

Group contrast (G1 / G4)	G1	G4	Ratio	95% CI
Income: poorest / richest	53.4	28.0	1.90	[1.77, 2.05]
Car-free: >30% / <10%	240.1	16.9	14.21	[13.54, 14.92]
Poverty: >60% / <20%	72.2	29.7	2.43	[2.28, 2.59]
Distance: core / fringe	124.0	2.4	50.75	[48.70, 52.89]
%Black: >60% / <10%	71.5	33.6	2.13	[1.99, 2.28]
%Hispanic: >60% / <10%	37.0	37.9	0.98	[0.92, 1.04]

45-minute budget, 2023, 52 metros outside New York, in thousands of jobs. Group 1 and Group 4 access (population-weighted means of block-group 2023 values, thousands) and their ratio with a delta-method 95% confidence interval, all jobs (C000); each row is the more transit-reliant group over its reference group, so a ratio above one favors the reliant group (ratios computed before rounding).

5.4 Within cities: who gained access?

In growth, the decade ran the other way, and on the corrected index the reversal is stark (Fig. 7). The urban fringe gained about 32% over the decade while the core gained little (+4%), and this distance gradient propagates into every demographic cut (Fig. 8). The richest income quartile gained about 31% while the poorest gained essentially nothing (−0.3%); the most car-owning neighborhoods gained about 19% while the most car-free gained essentially nothing; low-poverty neighborhoods gained about 22% while high-poverty neighborhoods gained barely a point. The starkest movement is racial: majority-White neighborhoods gained about 16% (majority-Hispanic about 11%), while majority-Black neighborhoods' job access by transit declined: by about 9% within 45 minutes and, for longer-reach trips, roughly 13% of the jobs reachable within 60 minutes. It is a decline concentrated in the handful of metros where majority-Black population concentrates. And the within-city gaps opened before the pandemic: in the distance-controlled framework of Table 3, the poorest–richest contrast had already reached roughly −20 log-points by 2019, most of its decade value. The gradient is not an artifact of the 45-minute budget: it holds at every threshold, steeper at 30 minutes (the richest quartile gained about +40% against the poorest's +5%) and shallower at 60 (about +26% against −2%). The majority-Black decline deepens with the budget: longer-reach access, which reaches more of the metro's jobs, eroded most. The uncorrected data had hidden it: the operators that entered the archive after 2014 disproportionately serve poorer and majority-Black neighborhoods, so archive adoption masked their losses. Because majority-Black population weight is concentrated in those few metros, the decline is a fact about them rather than an estimate of a uniform national process (its metro-clustered interval spans zero; Appendix A); the gains in richer, whiter, peripheral neighborhoods, by contrast, are metro-robust. The mechanism is the one the decomposition identifies (Section 5.1): the movement is chiefly jobs-side, reachable-job counts growing where jobs sprawled and falling where they hollowed, with the within-city accounting in Table 3. In percentage terms low-access peripheral neighborhoods converged toward high-access central ones; in absolute terms every change is small against the level differences (the fringe's +32% is under a thousand reachable jobs; the core's few-percent gain, several thousand). We report both readings and attach no welfare claim to either.

The growth gradient is not a composition of growing metros, nor of urban form alone. Comparing each neighborhood only to others in the same metro and the same five-kilometer distance band (Table 3), the income gradient survives essentially intact: the poorest quartile's access grew 20.7 log-points less than the richest's over the decade (95% CI $[-22.1, -19.2]$; metro-clustered $[-26.8, -14.5]$, robust at either level). The poverty gradient is similar. The racial gap attenuates by about two-thirds, location explaining most of it, but a majority-

Figure 7: Change in job access by transit across distance-to-core quartiles, 2014–2025

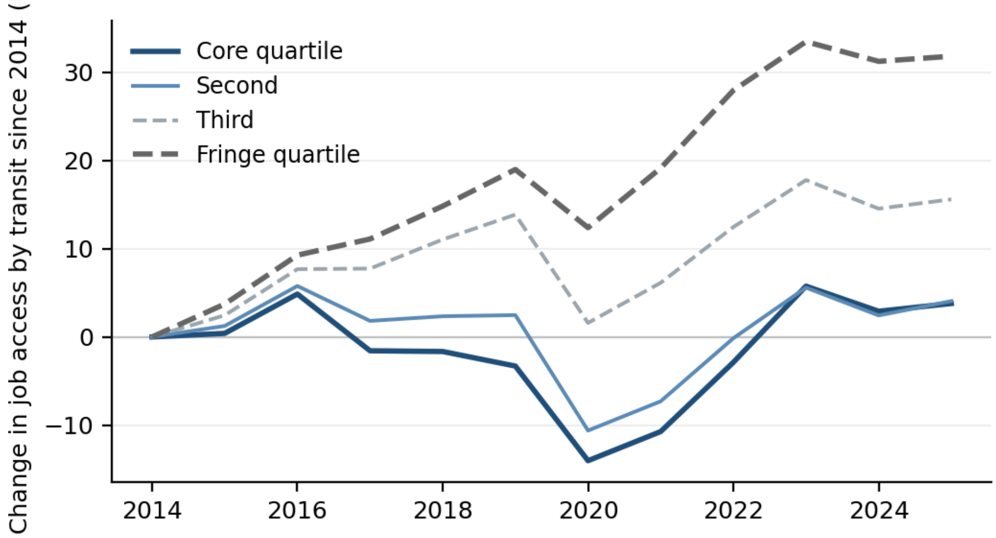

Within-block-group change in job access by transit relative to 2014, by distance-to-core quartile, annual, population-weighted, 52 metros outside New York, flagged cells excluded.

Figure 8: Decade change in job access by transit across neighborhood groups, corrected and uncorrected

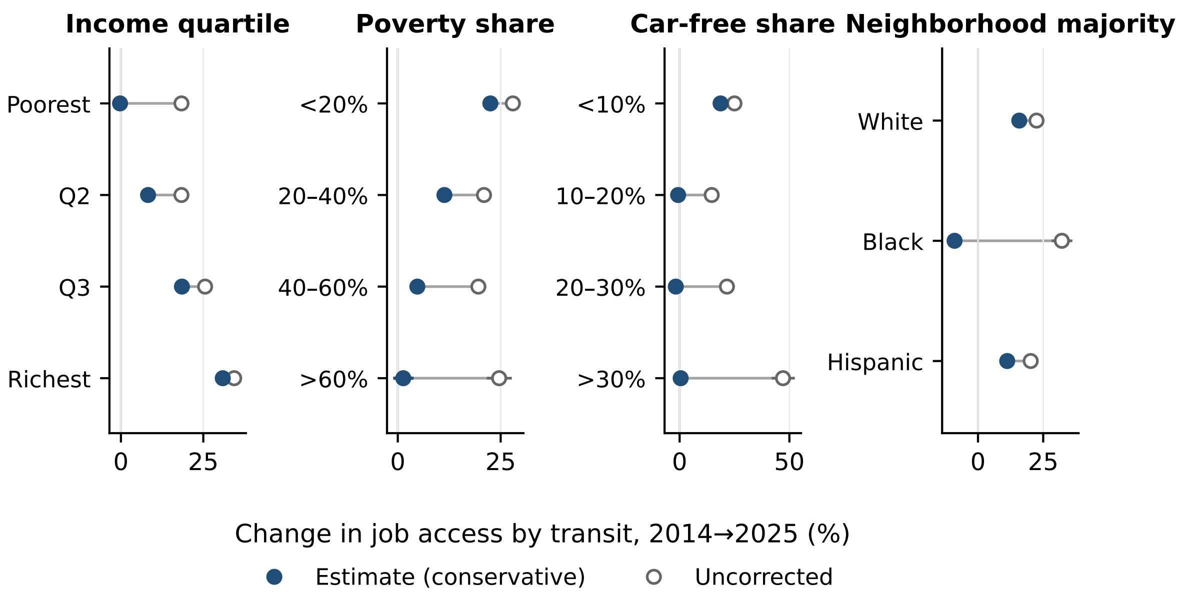

Change in job access by transit, 2014 to 2025, by neighborhood group. Filled markers: the estimate (conservative); open markers: the uncorrected value; whiskers are 95% block-group-clustered intervals. The distance between the two is the archive-adoption correction.

Black shortfall of about 11 log-points remains ($[-13.0, -8.3]$; under metro clustering it spans zero, Appendix A). Decomposing each contrast into its network and jobs components, with the archive-entry correction applied equally to the observed change and the network part so the columns approximately add, identifies the mechanism. For income, poverty, and car-reliance the gradient is on the jobs side: jobs-side losses of -19, -15, and -6 log-points against network components within a few points of zero. The racial contrast splits: about half is jobs-side (-5.1 log-points, adoption-free by construction), while the network part cannot be signed (it is -7 log-points if every archive entry is treated as adoption and $+7$ if every entry is genuine new service), so we do not lean on it. Within the same city and the same distance from the core, employment relocated toward richer neighborhoods over the decade; for the poorest-versus-richest contrast, transit service change played essentially no role. Because neighborhoods are classified at a fixed 2018–2022 vintage, this reads as "the places that are affluent today gained access"; whether jobs followed money or money followed access is a sorting question this comparison does not resolve. One version of the concern we can address is end-of-period sorting: classifying neighborhoods instead by their income at the start of the period (the 2009–13 ACS, crosswalked to 2020 block groups) leaves the gradient essentially intact. The quartile richest in 2009–13 gained about +31% over the decade and the poorest about zero, the same spread as under present-day classification, and the within-city distance-controlled gap is -16 log-points (metro-clustered $[-26, -5]$). The neighborhoods that were already better-off at the start of the period are the ones whose access grew; the pattern is not an artifact of who is rich today.

5.5 The transit-reliant margin

The introduction's motivating household is the one for which job access by transit is the opportunity set. Table 4 reads the panel on that margin directly, for the constrained transit-reliant strata defined in Section 3.1. All statistics weight block groups by their car-free households, so each row describes the typical car-free household in that stratum. In levels, the sorting documented cross-sectionally by the Modal Access Gap literature (Maharjan et al., 2024) is visible on the constrained margin too: the typical car-free household lives where transit reaches 131,000 jobs within 45 minutes, roughly 3.5 times the population-weighted average person's 38,000 (Table A.5), and even the poorest, most car-free strata reach 130–174,000, rising with the bar. Two disclosures attach. Membership is fixed at the 2018–2022 ACS vintage, so the strata are today's reliant neighborhoods traced backward. Classifying them on their 2009–13 attributes instead gives the same answer (constrained endpoints of -1.2 and -0.9 percent against $+3.0$ for the all-neighborhood car-free line), so the bypass is not an artifact of end-of-period sorting. And the strata concentrate in a few dense metros (Chicago, Los Angeles, and Miami together hold about 31% of the constrained car-free weight), with New York, the most extreme case, already analyzed separately.

Table 3: Within-city growth incidence, distance-controlled.

Contrast (within metro × distance band)	Observed	Network	Jobs
Poorest vs. richest income quartile	−20.7	−2.5	−18.8
	[−22.1, −19.2]	[−3.5, −1.5]	[−19.9, −17.7]
High- vs. low-poverty (>60% vs. <20%)	−17.9	−3.8	−15.4
	[−19.9, −15.9]	[−5.4, −2.2]	[−16.8, −14.1]
Majority-Black vs. majority-White	−10.7	−7.0	−5.1
	[−13.0, −8.3]	[−8.9, −5.0]	[−6.7, −3.4]
Car-free vs. car-owning (>30% vs. <10%)	−5.7	−0.5	−5.8
	[−7.6, −3.8]	[−2.4, +1.3]	[−6.7, −5.0]

Decade change in job access by transit across neighborhood contrasts, decomposed into network and jobs components, in log-points ×100. Each cell is the coefficient on the listed group contrast from a regression of the block group's decade change (2014–15 to 2024–25 window means, in logs, corrected for archive entry: each block group's cross-regime jumps at operator entries are removed) on group indicators, absorbing metro × 5-km-distance-band cells, population-weighted, robust standard errors (95% CIs in brackets); 70,057 block groups, the panel where both endpoint windows are observed (45 of 52 metros). Network and jobs columns use the per-block-group components of Section 5.1's decomposition construction, computed for this table's own sample, with the archive-entry jump removal applied equally to Observed and Network (the jobs component is adoption-free by construction), so rows approximately add; the small residual is the decomposition's ordinary approximation error. The majority-Black network entry is the corrected-basis estimate. Because each contrast differences within metro, metro-wide job growth cancels from the Jobs column.

Table 4: Job access by transit at the transit-reliant margin.

Stratum	Block groups	Car-free HH covered (%)	Access (000s)	Change 2014–25 (%)	Top-3 metro share (%)
All block groups	85,694	100	131	+2.5	29
Constrained: car-free ∩ below-200%-poverty					
>10% ∩ poor >40%	10,400	38	130	+1.1	31
>20% ∩ poor >50%	4,386	22	147	+2.9	34
>30% ∩ poor >50%	2,585	16	174	+2.8	38

52 metros outside New York, balanced panel, 2014–2025 pooled, flagged cells excluded. Access = jobs reachable within 45 minutes by transit (thousands, all jobs), weighted by car-free households (ACS 2018–2022: car-free from B25044, poverty share from C17002; membership fixed at that vintage). Change 2014–25 = the stratum's archive-entry-corrected within-block-group fixed-effects estimate for 2025 (Fig. 9). Top-3 metro share = the share of the stratum's car-free households living in its three largest metros. The reliance cut is a neighborhood-level double threshold on ACS marginals, not a household-level joint distribution (Section 7).

Over time, the decade's rise reaches this margin only faintly (Fig. 9). On the index, weighted by car-free households, access for the average car-free household rose just +2.5% over the decade (CI [+1.6, +3.3]), a small fraction of the +15% the average resident gained, and the constrained strata gained no more, ending between +1 and +3% above their 2014 level (block-group intervals that straddle or barely clear zero). The uncorrected series ends far higher (about +26%, Fig. 9); the gap between the two is precisely the operators entering the archive in these neighborhoods, the adoption the correction exists to remove (Section 4). The bypass is thus the conservative reading, and it is the one the within-city comparison below supports independently. These strata start high in levels, being central, but the decade's growth passed them by. The more telling comparison is relative and survives the urban-form control: comparing block groups within the same metro and the same five-kilometer distance band (the framework of Table 3), constrained neighborhoods' decade growth was 8 to 13 log-points lower than their same-distance neighbors' (metro-clustered intervals exclude zero). The shortfall is chiefly jobs-side (the jobs component alone is about 8 log-points, while the network component is small and spans zero across archive-entry treatments). It is employment relocating away from these neighborhoods rather than service withdrawal or distance to the core. This is the growth-incidence result of Section 5.4 on the population the motivation names: the population-weighted +15% is a rise for the average resident, far more than for the average car-free household, whose access barely moved.

Figure 9: Change in job access by transit, weighted by car-free households, 2014–2025

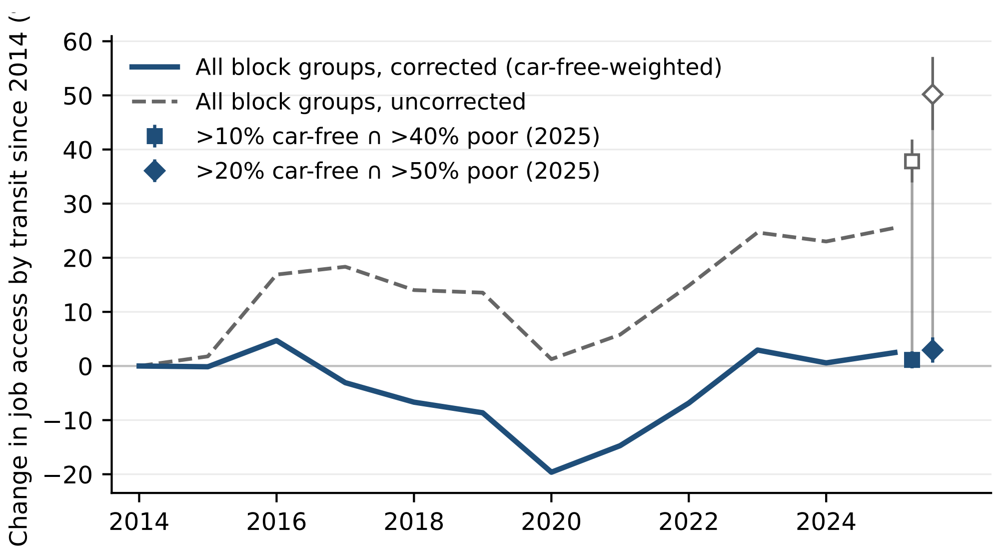

Change in job access by transit relative to 2014, annual, weighted by car-free households, all block groups, 52 metros outside New York, flagged cells excluded: the corrected series (solid) and the uncorrected series (dashed); the gap between them is archive adoption. At right, the 2025 decade change for the two broadest constrained strata (>10% car-free ∩ >40% poor, squares; >20% car-free ∩ >50% poor, diamonds): filled markers corrected, open markers uncorrected, whiskers 95% block-group-clustered.

5.6 Across metros: winners and losers

The national trajectory hides wide variation across metros that the method-consistent panel can rank (Fig. 10). The largest gains accrued to fast-growing Sun Belt and western metros (Austin, Las Vegas, Phoenix, Charleston, and Sacramento lead), while several large legacy systems lost ground: Miami, Springfield (Illinois), Chicago, Boston, and Kansas City. The correction matters here: Miami's raw series shows a +18% gain that is entirely operators entering the archive after 2015 (Fig. 3); once these entries are removed, the typical Miami neighborhood lost about 10%, the largest measurable decline. Chicago, whose degraded post-2020 archive we recovered from the operators' restored feeds, is now measurable too and declined about 7%, no longer the outlier its unaudited feed record once suggested. Flagged metro-quarters are excluded throughout (Appendix Table A.3); surfacing these artifacts is itself a caution: cross-metro comparisons on archived feeds require this auditing.

Figure 10: Decade change in job access by transit across metropolitan areas

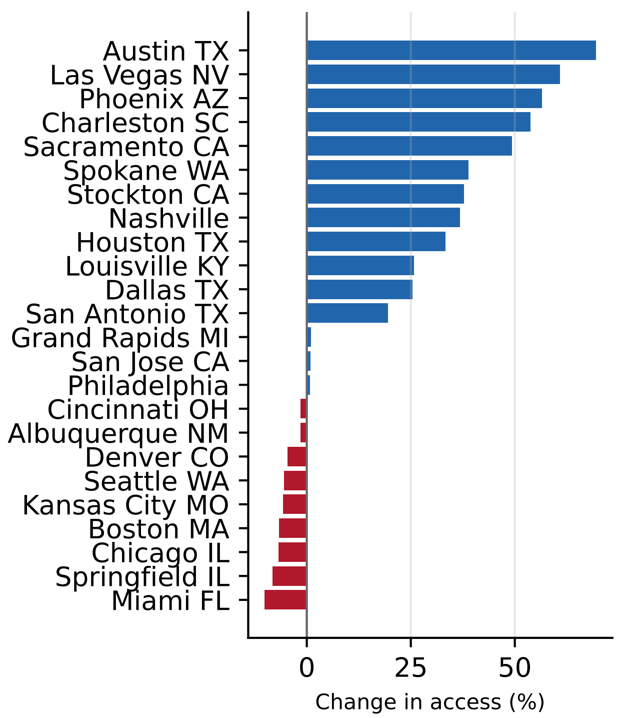

The typical neighborhood's change in job access by transit, 2014–15 to 2024–25, by metro (corrected, within block group, population-weighted; top and bottom 12 of the panel where both endpoint windows are observed, 45 of 52 metros—the seven without a clean window at one or both endpoints, and hence unranked, are Bremerton, Los Angeles, Oklahoma City, Providence, Redding, Salinas, and Washington).

5.7 The pandemic, and other episodes

COVID-19 is the episode the panel resolves best. The 2020 collapse and recovery are already visible in the trajectory (Fig. 6) and, mechanistically, in the decomposition: it is the network component, not the jobs component, that falls and fails to recover after 2020 (Fig. 5). This is the access counterpart of the service cuts that the COVID-transit literature has documented at the system and metro level (Kar et al., 2022). Agency-reported service data (the National Transit Database) corroborate the component's shape through the collapse and bound its incomplete recovery: all-day service miles returned faster than peak-hour job reach (Appendix A). The comparison is not circular. The core metros are never flagged by any source, the NTD scans included, so the exhibit does not select metros for agreeing with the agency data, and selection toward agreement could not produce the pre-2020 disagreement it shows. That the cuts depressed access, and hit socially vulnerable neighborhoods hardest, is established. Our contribution is to resolve the collapse and the recovery at the neighborhood level and quarter by quarter, which an annual series cannot, and to place it within the longer decade rather than treat it as the whole story. Figure 11 shows what that resolution buys: the collapse and the recovery as two change maps for Boston. The collapse (left) is broad but deepest in the middle-ring neighborhoods where frequency cuts truncated marginal service; the recovery (right) is its mirror image, except for the patches that stay red, the neighborhoods where service had not returned by 2025. The mapped area's median reachable-job count fell by about a fifth (−22%) in three quarters; an annual vintage would average through the entire movement. The remaining episodes are lighter texture: the 2015–19 bus-network-redesign wave, whose stated aim was reallocating service within the city (winners and losers that can cancel in the metro mean), and the post-2008 recession, which predates the quarterly panel and enters only as context.

6 New York: the exception

New York is not a large version of the other metros; it is a different object, and pooling it into the primary sample would let one city speak for the country. We therefore report it on its own; a dedicated, New-York-specific account of how its job access by transit and its racial distribution evolved is given by Zhang et al. (2026), and here we place New York in national relief. Four facts set it apart.

Scale. New York's neighborhoods reach about ten times the job access by transit of the typical metro elsewhere, about 396,000 jobs within 45 minutes against 38,000 outside New York (Table A.5, Appendix A). Its transit/car gap is correspondingly small: about 12× at

Figure 11: Block-group change in job access by transit, Boston, 2019Q4–2020Q3 and 2020Q3–2025Q4

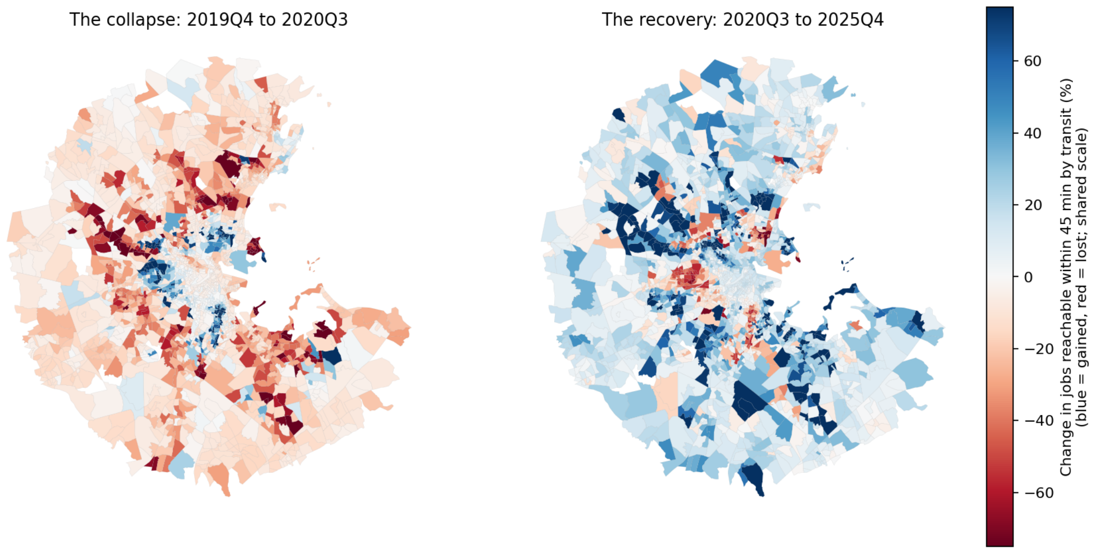

Percent change in jobs reachable within 45 minutes by transit, per block group, Boston: the collapse (2019Q4 to 2020Q3) and the recovery (2020Q3 to 2025Q4), shared diverging scale (central 30 km; base $\geq$ 100 jobs). Boston is chosen for its dense, unflagged archive coverage across the display window.

45 minutes, against 38× elsewhere, because transit in New York is genuinely competitive with the car. By every level measure New York sits an order of magnitude apart.

Trajectory and the pandemic. Composition-adjusted, New York's job access by transit rose about +11% over the decade against +25% outside it, both on the uncorrected index, the like-for-like comparison, since the archive-entry correction is poorly identified for a single metro. With no cross-metro variation to separate an archive-regime step from genuine service change, New York's dense history of feed reorganizations absorbs most of its measured growth into regime boundaries (a corrected New York index sits near zero, which we read as an artifact of the construction rather than of the city). On either footing, New York rose less than the metros outside it. Its access dipped below the 2014 baseline in 2020 (about −4%), as it did outside New York. But because New York holds so large a share of national job access by transit, pooling it into a national average would pull that average down and let one exceptional city dominate it, the reason we report the metros outside New York as our primary sample (Section 3.2).

The decomposition reverses. New York is the only place where the transit network's own contribution did not retreat: holding jobs fixed, its network adds about +1 log-point over the decade, against about −6 everywhere else (Fig. 12). The reading is intuitive: New York's network is mature and built out, with little room to expand and, after 2020, the institutional capacity to hold service where thinner systems cut it. Its rise is almost purely job growth (about +9 log-points; net job relocation is nil at +0.2), but the network-retreat finding that holds across the rest of the country is a non–New-York phenomenon.

Figure 12: Network, job-growth, and job-relocation components of access change: New York

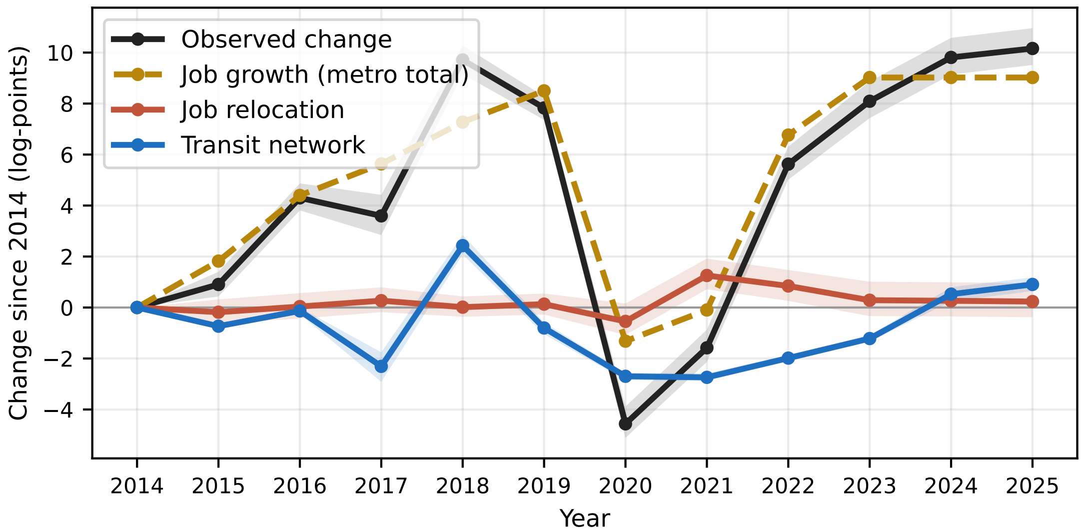

The decomposition of Fig. 5 for New York alone: transit network (jobs held at 2014), job growth, and job relocation components, within block group, annual, population-weighted, log-points, indexed to the 2014 annual average.

The gradients invert. Where access elsewhere is cleanly pro-poor and pro-minority, New York's distribution runs the other way (Table 5). Figure 13 shows the racial inversion on the ground, with the region's majority-Black neighborhoods sitting visibly off the access peaks (southeast Queens, the edges of central Brooklyn, Newark) rather than on them. The income gradient washes out: its richest and poorest neighborhoods reach essentially the same access (1.07×, not statistically distinct), because in New York both the affluent and the poor live centrally: a U-shape rather than a ladder. And the racial gradient reverses: majority-Black neighborhoods reach about a third fewer jobs than majority-White ones (0.66×), the opposite of the roughly 2.1× advantage outside New York. The reversal reflects where the city's Black neighborhoods sit, in eastern Brooklyn and the eastern Bronx, farther from the Manhattan job core. The car-reliance and distance gradients, by contrast, are even steeper than elsewhere (about 96× and 280×). Pooled into national figures these opposing patterns would partly cancel the gradients found everywhere else; separating New York keeps both stories clean.

Figure 13: Job access by transit per block group, New York, 2023Q4, majority-Black block groups outlined

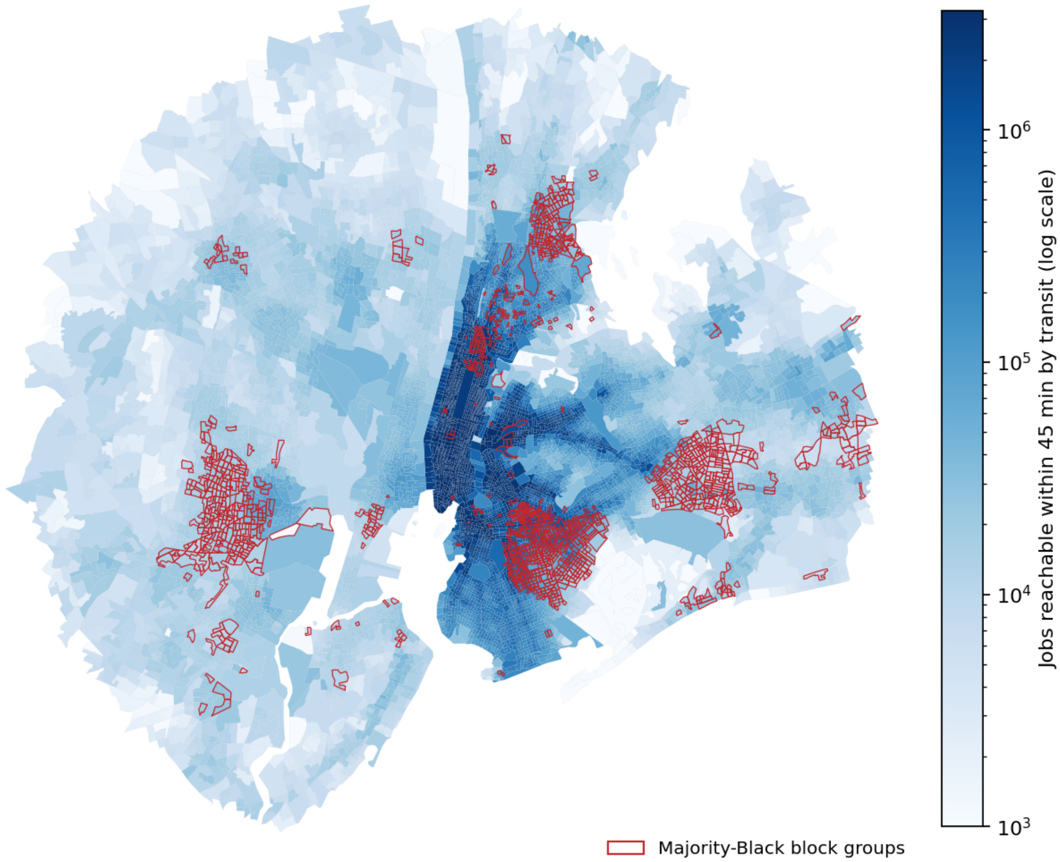

Jobs reachable within 45 minutes by transit, per block group, 2023Q4 (log scale, central 35 km), with the region's 1,513 majority-Black block groups outlined.

Table 5: New York: job access by transit across neighborhood groups.

Group contrast (G1 / G4)	G1	G4	Ratio
Income: poorest / richest	466.7	436.7	1.07
Car-free: >30% / <10%	992.0	10.4	95.70
Poverty: >60% / <20%	579.1	319.2	1.81
Distance: core / fringe	1,283.5	4.6	279.90
%Black: >60% / <10%	259.2	393.1	0.66
%Hispanic: >60% / <10%	443.6	359.8	1.23

45-minute budget, pooled 2014–2025, in thousands of jobs; population-weighted means of block-group values, all jobs (C000). Groups are defined as in Table 2.

7 Discussion and Conclusion

We read these facts as bounded measurement. The central one, that the decade's rise in job access by transit, outside New York, is a jobs-side story: peak-period scheduled service contributed roughly nothing to it before 2020 and subtracted from it after, disciplines an intuitive prior that rising access reflects transit investment. The network component's post-2020 level is bounded rather than point-identified: its metro-clustered interval spans zero, and agency-reported all-day service recovered faster than peak-hour job reach (Appendix A). Its near-zero contribution through 2019 and its 2020 collapse are not in doubt. The reading is descriptive: we document where the jobs a rider can reach came from, not whether reaching them changed anyone's employment.

Several limitations bound the reading, and we state them plainly. First, GTFS encodes scheduled rather than delivered service, so we measure timetabled access, not the reliability a rider experiences. Second, the no-substitution rule leaves coverage holes (thin 2013–2015 and a 2020–2023 archival trough) which shrink the sample but, because we drop rather than impute, do not bias the cells that remain. The archive's own expansion over 2015–2016 is handled by the archive-entry correction that all trend estimates carry (Section 4); for the subgroups where operators entered the archive late, majority-Black and car-free neighborhoods above all, the corrected and uncorrected readings differ widely, so we show both (Figs. 8, 9) and lead on the corrected one. Third, the car baseline is free-flow on a present-day street network, so the gap's level is an upper bound (Section 3.1); the transit-access trajectory that carries the temporal narrative is insensitive to it. Fourth, cumulative-opportunity access counts reachable jobs and is silent on competition for them, on within-window frequency, and on transfer disutility, and on fares, reliability, crowding, and safety, so what we measure is schedule feasibility. It is potential rather than effective access, and the competition omission bites hardest for the low-wage workers most dependent on transit (Shen, 1998). Fifth, the panel prices the morning peak only: accessibility varies strongly within the day (Boisjoly and El-Geneidy, 2016), off-peak and weekend access matter most for non-standard work schedules. As the agency-service comparison of Fig. A.1 (Appendix A) shows, the post-2020 recovery looks different off-peak than at the peak our measure holds fixed. Sixth, because LODES records a job at its employer's establishment rather than where the work is performed, the rise of remote work lowers the welfare relevance of physical access for telework-able jobs after 2020 (Dingel and Neiman, 2020). Reweighting the jobs side to physically-present, non-remote employment (using industry teleworkability shares) leaves the relocation finding intact (non-remote jobs relocated toward transit-served places as much as all jobs did), so the jobs-side story is not an artifact of remote-able jobs recorded at downtown establishments. Seventh, the transit-reliant cut is a neighborhood-level (car-light $\cap$ low-income) intersection, not a household-level one (Section 3.1).

Implications for measurement and practice. Three uses follow from the measurement alone. First, service supplied and job reach are different recovery metrics: in the clean-archive metros, all-day vehicle revenue miles ended the decade about 7% above their 2014 level while the network's contribution to peak job reach ended about six log-points below zero (Appendix A), so a dashboard tracking service hours alone overstates the recovery of work access. Peak job reach is computable from an agency's own published schedules. Second, access statistics weighted by the average resident overstate what transit-dependent households experienced: weighted by car-free households, the decade's rise falls from +15% to +2.5% (Section 5.5), so equity reporting should weight by the population the claim is about. Third, archived schedules support trend research only after auditing: all but three of the metro-quarters that external agency data ultimately flagged had passed every internal consistency check (Section 3.2), so trend work on archived GTFS should corroborate the archive against an independent service series.

We have built the first consistent, neighborhood-level, quarterly panel of U.S. job access by transit on a single differenceable methodology, 2014–2025, and used it to document how that access evolved. Composition-adjusted and corrected for the archive's own expansion, the typical neighborhood outside New York saw its job access by transit rise about 15% over the decade, with a COVID-19 collapse that briefly erased the entire gain and pushed access below its 2014 level, then a recovery that carried it above its pre-pandemic peak. The panel resolves that arc at the neighborhood level, quarter by quarter, as annual series cannot. Decomposing that change, an exercise only such a longitudinally comparable panel permits, shows the rise rested entirely on jobs: metro job growth alone would have delivered it. The transit network's own post-pandemic retreat took back a substantial part, and job relocation toward transit-served places recovered most of that loss. The added reach is the work of employment growth; scheduled peak-period service ended the decade subtracting from it. Access levels favor poorer, more central, and car-free neighborhoods throughout, but the decade's growth ran the other way, landing in the peripheral, car-owning, higher-income neighborhoods where the jobs sprawled, while the poorest gained essentially nothing and majority-Black neighborhoods ended the decade below their 2014 access. The transit/car gap remains an order of magnitude or more, the spatial-mismatch penalty facing households without a car. New York, set aside from these results for the reasons given, is documented on its own (Section 6).

What the panel sets up, finally, is a consistent, longitudinally comparable measure of job access by transit: an input the causal literature has so far assembled from heterogeneous, methods-shifting sources.

CRediT authorship contribution statement

Sierra Arnold: Conceptualization, Data curation, Formal analysis, Investigation, Methodology, Software, Validation, Visualization, Writing – original draft, Writing – review & editing.

References

Abadie, A., Athey, S., Imbens, G.W., Wooldridge, J.M., 2023. When should you adjust standard errors for clustering? The Quarterly Journal of Economics 138, 1–35. doi:10.1093/qje/qjac038.

Bastiaanssen, J., Johnson, D., Lucas, K., 2020. Does transport help people to gain employment? A systematic review and meta-analysis of the empirical evidence. Transport Reviews 40, 607–628.

Boisjoly, G., El-Geneidy, A., 2016. Daily fluctuations in transit and job availability: A comparative assessment of time-sensitive accessibility measures. Journal of Transport Geography 52, 73–81.

Conway, M.W., Byrd, A., van Eggermond, M., 2018. Accounting for uncertainty and variation in accessibility metrics for public transport sketch planning. Journal of Transport and Land Use 11, 541–558. doi:10.5198/jtlu.2018.1074.

Dingel, J.I., Neiman, B., 2020. How many jobs can be done at home? Journal of Public Economics 189, 104235. doi:10.1016/j.jpubeco.2020.104235.

Ermagun, A., Witlox, F., 2024. Transit access effectiveness in American metropolitan areas. Journal of Transport Geography 116, 103853.

Fink, C., Klumpenhouwer, W., Saraiva, M., Pereira, R.H.M., Tenkanen, H., 2022. r5py: Rapid Realistic Routing with R5 in Python. doi:10.5281/zenodo.7060437. [software].

Foth, N., Manaugh, K., El-Geneidy, A., 2013. Towards equitable transit: examining transit accessibility and social need in Toronto, Canada, 1996–2006. Journal of Transport Geography 29, 1–10.

Geurs, K.T., van Wee, B., 2004. Accessibility evaluation of land-use and transport strategies: review and research directions. Journal of Transport Geography 12, 127–140.

Glaeser, E.L., Kahn, M.E., 2001. Decentralized employment and the transformation of the American city. Brookings-Wharton Papers on Urban Affairs 2001, 1–63.

Gobillon, L., Selod, H., Zenou, Y., 2007. The mechanisms of spatial mismatch. Urban Studies 44, 2401–2427.

Grengs, J., 2010. Job accessibility and the modal mismatch in Detroit. Journal of Transport Geography 18, 42–54.

Handy, S.L., Niemeier, D.A., 1997. Measuring accessibility: an exploration of issues and alternatives. Environment and Planning A 29, 1175–1194.

Hansen, W.G., 1959. How accessibility shapes land use. Journal of the American Institute of Planners 25, 73–76.

Ihlanfeldt, K.R., Sjoquist, D.L., 1998. The spatial mismatch hypothesis: a review of recent studies and their implications for welfare reform. Housing Policy Debate 9, 849–892.

Janatabadi, F., Maharjan, S., Ermagun, A., 2023. A spatiotemporal disparity of transit and automobile access gap and its impact on transit use. Environment and Planning B 50, 1858–1878.

Kain, J.F., 1968. Housing segregation, negro employment, and metropolitan decentralization. Quarterly Journal of Economics 82, 175–197.

Kar, A., Carrel, A.L., Miller, H.J., Le, H.T.K., 2022. Public transit cuts during COVID-19 compound social vulnerability in 22 US cities. Transportation Research Part D 110, 103435.

Klumpenhouwer, W., Allen, J., Li, L., Liu, R., Robinson, M., Da Silva, D., Farber, S., Karner, A., Rowangould, D., Shalaby, A., Buchanan, M., Higashide, S., 2021. A comprehensive transit accessibility and equity dashboard. Findings doi:10.32866/001c.25224.

Kneebone, E., 2009. Job Sprawl Revisited: The Changing Geography of Metropolitan Employment. Technical Report. Metropolitan Policy Program, Brookings Institution. Washington, DC.

Lahoorpoor, B., Rayaprolu, H., Wu, H., Levinson, D.M., 2022. Access-oriented design? Disentangling the effect of land use and transport network on accessibility. Transportation Research Interdisciplinary Perspectives 13, 100536. doi:10.1016/j.trip.2021.100536.

Liu, L., Miller, H.J., Scheff, J., 2020. The impacts of COVID-19 pandemic on public transit demand in the United States. PLOS ONE 15, e0242476.

Maharjan, S., Janatabadi, F., Ermagun, A., 2024. Spatial inequity of transit and automobile access gap across America for underserved population. Transportation Research Record 2678, 674–690.

Ohio Supercomputer Center, 1987. Ohio Supercomputer Center. URL: https://ror.org/01apna436. Columbus, OH: Ohio Supercomputer Center.

Owen, A., Levinson, D.M., 2015. Modeling the commute mode share of transit using continuous accessibility to jobs. Transportation Research Part A 74, 110–122.

Owen, A., Liu, S.S., Lind, E.M., 2025. Access Across America: Transit 2024. Technical Report CTS 25-18. Center for Transportation Studies, University of Minnesota.

Owen, A., Murphy, B., 2020. Access Across America: Transit 2019. Technical Report CTS 20-09. Center for Transportation Studies, University of Minnesota.

Owen, A., Murphy, B., Levinson, D.M., 2016. Access Across America: Transit 2015. Technical Report CTS 16-09. Center for Transportation Studies, University of Minnesota.

Pereira, R.H.M., Saraiva, M., Herszenhut, D., Braga, C.K.V., Conway, M.W., 2021. r5r: rapid realistic routing on multimodal transport networks with R5 in R. Findings doi:10.32866/001c.21262.

Pereira, R.H.M., Schwanen, T., Banister, D., 2017. Distributive justice and equity in transportation. Transport Reviews 37, 170–191.

Pew Research Center, 2024. 1 in 10 Americans rarely or never drive a car. URL: https://www.pewresearch.org/short-reads/2024/11/14/1-in-10-americans-rarely-or-never-drive-a-car/. Short Read, November 14 (nationally representative survey of 5,410 U.S. adults, August 12–18, 2024).

Rosik, P., Mazur, M., Stepniak, M., Komornicki, T., Goliszek, S., Duma, P., Churski, P., 2025. What drives potential accessibility decomposition? Temporal and spatial variability of the impact of infrastructure and population components in France, Spain, and Poland in the years 1960–2020. Journal of Transport and Land Use 18. doi:10.5198/jtlu.2025.2508.

Sanchez, T.W., 1999. The connection between public transit and employment: the cases of Portland and Atlanta. Journal of the American Planning Association 65, 284–296.

Santana Palacios, M., El-Geneidy, A., 2022. Cumulative versus gravity-based accessibility measures: which one to use? Findings doi:10.32866/001c.32444.

Schroeder, J., Van Riper, D., Manson, S., Cooper, G., Krause, Z., Kugler, T., Zhu, T., Ruggles, S., 2026. IPUMS National Historical Geographic Information System: Version 21.0. doi:10.18128/D050.V21.0. [dataset]. Minneapolis, MN: IPUMS.

Shen, Q., 1998. Location characteristics of inner-city neighborhoods and employment accessibility of low-wage workers. Environment and Planning B: Planning and Design 25, 345–365. doi:10.1068/b250345.

Shorrocks, A.F., 2013. Decomposition procedures for distributional analysis: a unified framework based on the Shapley value. The Journal of Economic Inequality 11, 99–126. doi:10.1007/s10888-011-9214-z.

Stepniak, M., Rosik, P., 2018. The role of transport and population components in change in accessibility: the influence of the distance decay parameter. Networks and Spatial Economics 18, 291–312.

Tomer, A., Kneebone, E., Puentes, R., Berube, A., 2011. Missed Opportunity: Transit and Jobs in Metropolitan America. Technical Report. Brookings Institution Metropolitan Policy Program.

Tyndall, J., 2017. Waiting for the R train: public transportation and employment. Urban Studies 54, 520–537. doi:10.1177/0042098015594079.

U.S. Census Bureau, 2023. American Community Survey, 1-Year Estimates. Table B25044, Tenure by Vehicles Available.

Wessel, N., Allen, J., Farber, S., 2017. Constructing a routable retrospective transit timetable from a real-time vehicle location feed and GTFS. Journal of Transport Geography 62, 92–97.

Yeganeh, A.J., Hall, R.P., Pearce, A.R., Hankey, S., 2018. A social equity analysis of the U.S. public transportation system based on job accessibility. Journal of Transport and Land Use 11, 1039–1056. doi:10.5198/jtlu.2018.1370.

Zhang, Y., Xu, F., Talen, E., González, M.C., Li, Y., 2026. Quantifying racial inequality in transit access across New York City. PNAS Nexus 5, pgag025. doi:10.1093/pnasnexus/pgag025.

A Supporting exhibits

Figures

Figure A.1: Transit-network component of access and NTD vehicle revenue miles, 2014–2025

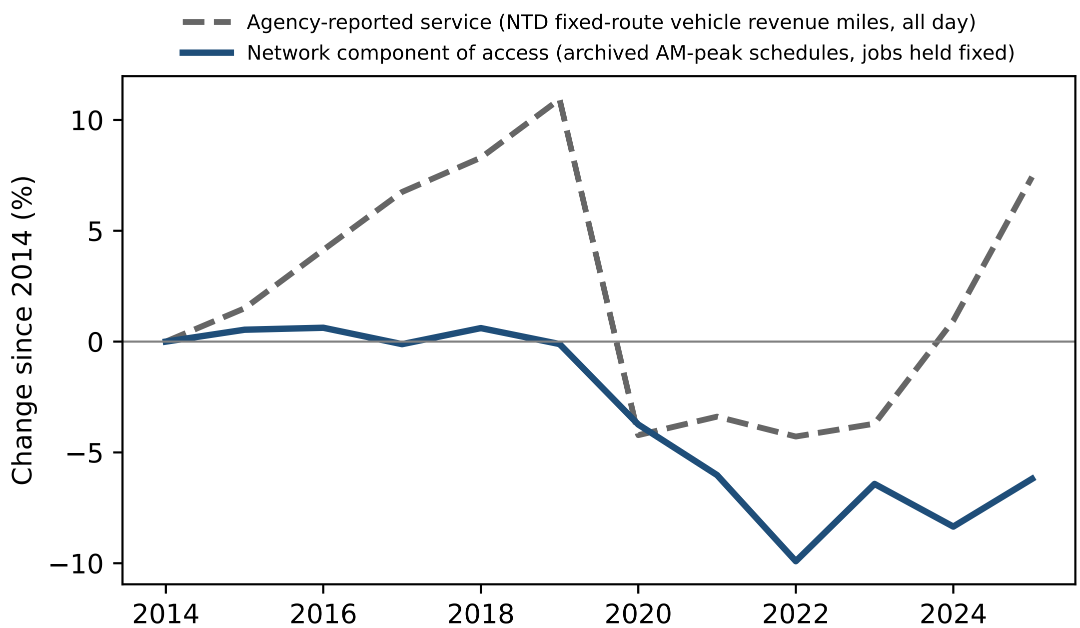

Change since 2014 in agency-reported fixed-route vehicle revenue miles (VRM, from the National Transit Database, the standard measure of service supplied, at all hours; dashed) and in our transit-network component (archived Wednesday-AM-peak schedules, jobs held fixed; solid), pooled over the clean-archive core (19 metros), all with complete 2014–2025 NTD coverage. VRM stands about +11% above its 2014 level by 2019 and about +7% by 2025; the network component is within a point of zero by 2019 and about −6 by 2025. The axis is in percent, the natural unit for the miles comparison; the decomposition exhibits are in log-points (Section 5.1).

Figure A.2: Decomposition component paths by metro archive status

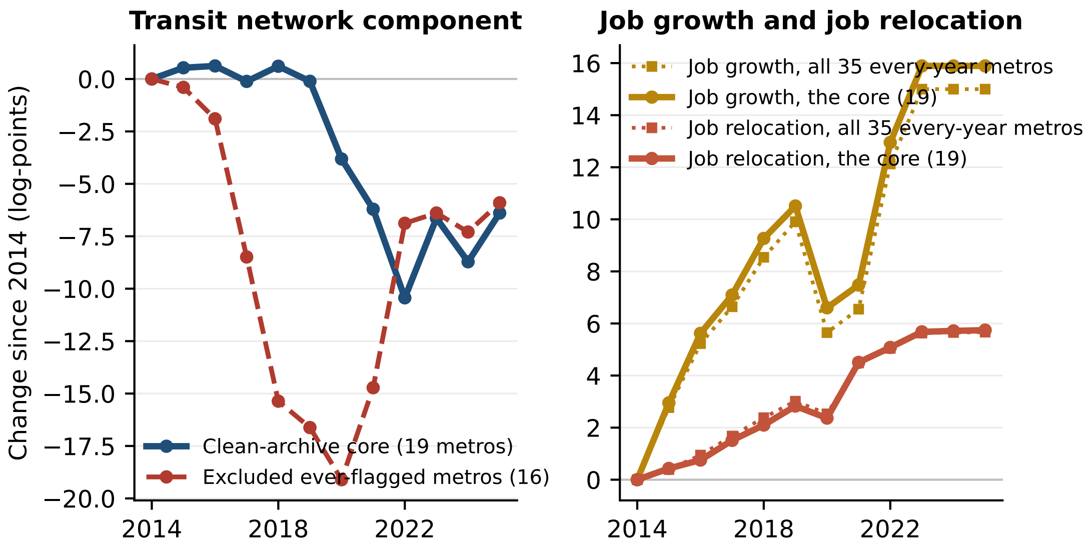

Left: the transit-network component on the clean-archive core (19 metros, solid) and on the 16 every-year metros excluded for flagged quarters (dashed); corrected, within block group, population-weighted, log-points vs the 2014 annual average. Right: the job-growth and job-relocation components on all 35 every-year metros (dotted) and on the core (solid), same construction. Flags mark metro-eras, so an affected metro is excluded whole rather than quarter by quarter (Section 3.2).

Figure A.3: Access levels against published Access Across America values

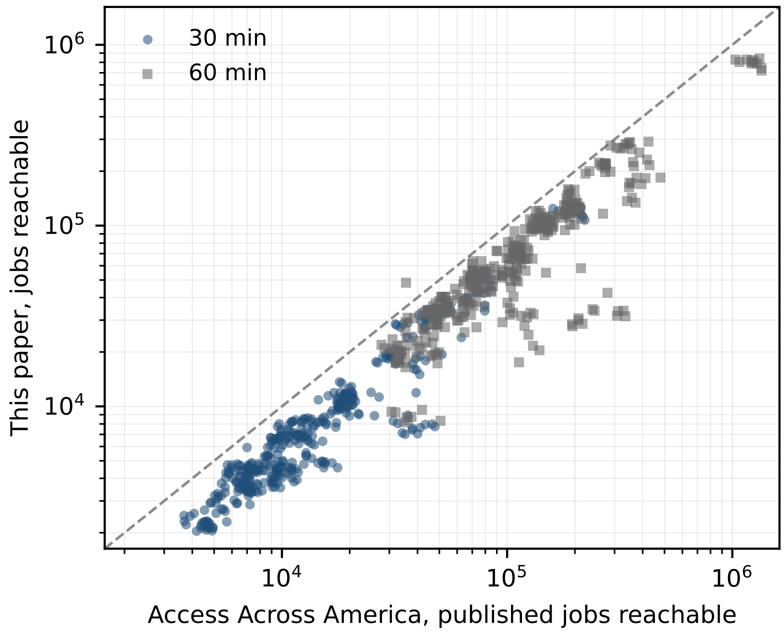

Our population-weighted metro-year access against the worker-weighted values published in the *Access Across America: Transit* reports (Owen et al., 2016, 2025), at the 30- and 60-minute thresholds, log scales; the dashed line is equality. Across the 390 overlapping metro-years (43 metros, 2014–2024) the correlation in logs is 0.95 at 30 minutes and 0.88 at 60, and our levels sit at a median 0.57 and 0.64 of AAA's (interquartile ranges 0.47–0.64 and 0.54–0.72). Weighting and geography differ between the two series (Section 4); flagged cells are excluded on our side (Section 3.2).

Tables

Table A.1: The 52 panel metros and quarters present (of 48, 2014–2025).

Metro	Qtrs	Clean	Metro	Qtrs	Clean
Albany, NY	48	47	Milwaukee, WI	48	46
Albuquerque, NM	48	48	Nashville, TN	39	39
Anchorage, AK	46	45	Oklahoma City, OK	38	35
Atlanta, GA	46	38	Oxnard, CA	44	44
Austin, TX	46	46	Philadelphia, PA-NJ	48	45
Boston, MA-NH	48	46	Phoenix, AZ	48	48
Bremerton, WA	48	28	Portland, OR-WA	47	47
Champaign, IL	47	47	Providence, RI-MA	39	33
Charleston, SC	48	48	Redding, CA	39	39
Chicago, IL-IN	44	44	Reno, NV	48	47
Cincinnati, OH-KY-IN	48	48	Rochester, NY	47	47
Cleveland, OH	46	44	Sacramento, CA	47	47
Columbus, OH	42	38	Salem, OR	46	35
Dallas, TX	48	42	Salinas, CA	37	37
Denver, CO	48	48	Salt Lake, UT	44	43
Duluth, MN-WI	48	46	San Antonio, TX	45	45
Eugene, OR	42	41	San Diego, CA	46	44
Grand Rapids, MI	44	44	San Jose, CA	46	44
Houston, TX	48	47	Santa Cruz, CA	40	38
Indianapolis, IN	43	37	Santa Rosa, CA	48	48
Kansas City, MO-KS	45	45	Seattle, WA	45	45
Las Vegas, NV	43	40	Spokane, WA	48	48
Los Angeles, CA	39	27	Springfield, IL	42	41
Louisville, KY-IN	37	35	Springfield, MA	46	32
Madison, WI	48	48	Stockton, CA	46	46
Miami, FL	38	37	Washington, DC*	43	27

Quarters present (of 48) count flagged cells as present; "Clean" counts the quarters that enter value figures (unflagged). Metros contributing unflagged cells by year: 45 (2014), 48 (2015), 50–52 (2016–2023), 48 (2024), 44 (2025). *Washington, DC: bus operators only, no archived WMATA feed (Section 3.3).

Table A.2: Each exhibit's sample.

Exhibit	Sample	Metros	Block groups	Weights
Trajectory (Fig. 6); group paths (Figs. 7, 8)	The panel	52	85,694	pop.
Levels (Tables 2, A.5)	The panel	52	85,694	pop.
Transit-reliant (Table 4, Fig. 9)	The panel	52	85,694	car-free hh.
Within-city contrasts (Table 3)	Panel, both endpoint windows	45	70,057	pop.
Metro winners (Fig. 10)	Panel, both endpoint windows	45	—	pop.
Decomposition (Fig. 5); NTD check (Fig. A.1)	Clean-archive core	19	28,147	pop.
Map (Fig. 11)	Illustrative metro	1	—	—
New York (Section 6)	New York	1	—	pop.
AAA validation (Fig. A.3)	Panel–AAA overlap	43	—	pop.
Commute-share validation (Fig. 4)	The panel + New York	51	—	—

The samples: *the panel* (Section 3.2; balanced 52 metros outside New York) and *the clean-archive core* (of the panel, the 35 metros observed unflagged in every year, and of those the 19 with no flagged quarter across the decade). The core's members: Albuquerque, Austin, Champaign–Urbana, Charleston, Chicago, Cincinnati, Denver, Kansas City, Madison, Oxnard–Thousand Oaks, Phoenix, Portland, Rochester, Sacramento, San Antonio, Santa Rosa, Seattle, Spokane, and Stockton, holding 33% of the panel's population. Robustness exhibits (thresholds, telework, beginning-of-period attributes) use the panel; Fig. A.2 compares the two.

Table A.3: Archive-integrity flag classes.

Class	Detection rule	Cells
Partial outage	an operator present in adjacent quarters is absent from the metro-quarter's archived feed set	73
Thin feed	a metro-quarter's archived schedules carry implausibly little service against the metro's own series	6
Archive scan	level spikes from duplicated feed bundles, and sustained depressions from archive dropout windows, in the full-panel integrity scan	132
NTD divergence	archived-schedule service diverges from agency-reported vehicle revenue miles; corroborated metro-eras (Appendix A)	72

Cells are flagged metro-quarters per class; classes overlap, and some flagged cells fall outside the balanced panel, so the distinct flagged total inside the panel is 136 of 2,330 metro-quarters (5.8%). Flagged cells are excluded from every value figure and count as present for coverage (Section 3.2). Of the 72 NTD-corroborated cells, 69 had passed every internal check.

Table A.4: Representativeness: the 52-metro panel versus all U.S. metros.

Metro-level median	Panel (52 metros)	All U.S. metros (392)
Population	1.6M	0.24M
Transit commute share (%)	1.7	0.7
Median household income ($000)	80.5	67.5
Zero-vehicle households (%)	6.0	6.0
Non-Hispanic White (%)	63.1	70.8
Non-Hispanic Black (%)	6.5	6.3
Hispanic (%)	13.4	8.1
Non-Hispanic Asian (%)	4.5	2.0

Medians across metropolitan statistical areas, ACS 2018–2022 (tables B01003, B08301, B19013, B03002, B25044); New York excluded.

Table A.5: Jobs reachable by time budget.

Panel A. Balanced 52-metro panel outside New York, 2023

Time budget	Transit	Car	Gap
15 min	1.9	164.1	86×
30 min	10.9	780.5	71×
45 min	37.6	1,433.7	38×
60 min	86.9	1,868.2	21×

Panel B. New York, pooled 2014–2025

Time budget	Transit	Car	Gap
15 min	10.9	505.5	46×
30 min	119.1	2,678.8	22×
45 min	396.4	4,813.6	12×
60 min	800.4	6,445.4	8×

Population-weighted mean jobs reachable within each budget, in thousands, all jobs (C000); “Gap” = car/transit. Panel A excludes flagged cells; Panel B pools 2014–2025, over which New York’s levels are stable (the ex–New-York panel uses 2023 per the levels convention). Car is free-flow, so the gap is an upper bound on the car advantage (Section 3.1).

Table A.6: Correlation among access constructs.

	Cumulative	Gravity	Logsum
Cumulative (45 min)	1.00	0.90 (0.90)	0.92 (0.91)
Gravity	0.90 (0.90)	1.00	1.00 (1.00)
Logsum	0.92 (0.91)	1.00 (1.00)	1.00

Pearson correlations of block-group access scores in logs, from the routed 2019Q4 travel-time matrices (a dense, pre-pandemic quarter), $n = 85{,}694$ block groups, 52 metros outside New York; pooled, with the median within-metro correlation in parentheses. Constructs: the paper's *cumulative* measure (jobs reachable within 45 minutes); a *gravity* measure with negative-exponential impedance, half-life 45 minutes ($\sum_j \text{jobs}_j\, e^{-(\ln 2/45)\, t_{ij}}$); and a destination-choice *logsum* ($\ln \sum_j \text{jobs}_j\, e^{-t_{ij}/45}$). The two smooth-impedance parameters are deliberately distinct (identical impedance would make the logsum the log of the gravity measure by construction); both are calibrated to the paper's 45-minute budget, following the common-budget logic of Santana Palacios and El-Geneidy (2022).

Table A.7: Main statistics under block-group and metro clustering.

Statistic	Estimate	BG-clustered	Metro-spec est.	Metro-clustered
Trajectory, 2025 (corrected)	+14.1	[+13.3, +14.8]	+14.9	[+6.1, +23.8]
at 2019Q4	+8.9	[+8.1, +9.7]	+5.2	[−9.2, +19.6]
Decomposition: job growth	+15.9	—	+15.0	[+11.8, +18.2]
job relocation	+5.7	[+5.1, +6.4]	+5.3	[+2.8, +7.8]
transit network	−6.4	[−7.0, −5.9]	−4.9	[−17.6, +7.9]
Growth, richest quartile	+30.9	[+28.8, +33.0]	+30.9	[+17.6, +45.7]
Growth, majority-Black (45 min)	−9.0	[−11.4, −6.5]	−9.0	[−18.2, +1.2]
Within-city: poorest−richest	−20.7	[−22.1, −19.2]	−20.7	[−26.8, −14.5]
majority-Black−White	−10.7	[−13.0, −8.3]	−9.4	[−19.4, +0.6]

Bracketed ranges are 95% confidence intervals. Trajectory and group growth: archive-entry-corrected within-BG FE (trajectory in log-points relative to the 2014 base, group growth in percent). Decomposition: 2025 vs. the 2014 average, log-points; job growth is an accounting identity at the block-group level, so only its metro-level variation is reported. Within-city contrasts: metro × 5-km cells absorbed, log-points. Point estimates are the harmonized figures of Section 5.1–5.4; the metro-clustered specification uses the 2014Q1 base on the full balanced panel, and its own point estimates are shown beside its intervals. The two sets of points differ through the base, correction, and (for the decomposition rows) sample treatment without altering the metro-level significance. Metro clustering: ≈ 52 clusters (45 for the within-city rows).

Table A.8: Reconciliation of the trajectory and decomposition estimands.

	2025 endpoint
(1) Trajectory of record (52 metros, quarterly, regime-FE correction)	+14.1
(2) Same estimator on the canonical every-year subset	+14.7
(3) Same estimator on its never-flagged metros	+16.1
(4) Decomposition observed (annual, per-BG jump removal, same sample)	+15.2
Coverage sample step, (1)→(2)	+0.6
Archive-quality sample step, (2)→(3)	+1.4
Construction step, (3)→(4)	−0.9

The trajectory estimand and the decomposition estimand, reconciled on a common basis (Section 5.1), 2025 vs. 2014 in log-points: the two sample steps are the canonical every-year restriction and then the clean-archive (never-flagged) restriction; the construction step is annual endpoint means with the per-block-group correction form against quarterly regime fixed effects.