Off the Beaten Tract: Constructing a New Neighborhood Geography Using Revealed Preference

Evan Mast and Alaina Barca*

March 23, 2026

Abstract

We construct a new neighborhood geography using a revealed preference intuition: If people disproportionately move within neighborhoods, their boundaries can be backed out from migration flows. Our "districts," which consist of about nine census tracts each, correspond to recognizable local areas, as their boundaries align with physical barriers, sharp demographic changes, and local government borders. To illustrate applications, we first show that tract-level analyses of neighborhood sorting miss important broader patterns. Second, aggregating tract-level intergenerational mobility estimates to the district level increases precision threefold while introducing little aggregation bias, suggesting that districts could help solve recurring statistical power problems in research on neighborhood opportunity.

JEL Classification Codes: R23, R00

Keywords: Neighborhood definition, residential mobility, residential sorting

*The views expressed herein are solely those of the authors and do not necessarily represent the views of the Federal Reserve Bank of Philadelphia, the Federal Reserve System, or the U.S. Census Bureau. Any errors or omissions are the responsibility of the authors. The Census Bureau's Disclosure Review Board and Disclosure Avoidance Officers have reviewed this information product for unauthorized disclosure of confidential information and have approved the disclosure avoidance practices applied to this release. This research was performed at a Federal Statistical Research Data Center under FSRDC Project 0185 (CBDRB-FY25-0185). Thanks to John Voorheis for providing invaluable help with navigating the data environment and disclosure process. Thanks to Dionissi Aliprantis, Kyle Butts, Salim Furth, George Galster, John Logan, Jeff Lin, David Phillips, Davin Reed, Anna Tranfaglia, Cora Wigger, Anna Ziff, and Oren Ziv for providing helpful comments.

Introduction

Neighborhood geographies are fundamental to the social sciences. They are used to measure the quality of life and economic opportunity at a location, to define the unit of residential choice, and to capture spatial variation in local conditions or demographic composition. However, constructing these geographies is a difficult and unsolved problem. How large is a neighborhood? Do their sizes vary? How does one identify borders between similar areas?

One reason these questions are difficult to answer is that we do not know what area determines a given local attribute. Qualitative work suggests that different features are determined at different scales (Suttles 1972; Galster 2001). A street block may capture immediate neighbors or the areas in which children play. A broader scale, such as the "districts" of roughly 50,000 people discussed in Jacobs (1961)'s seminal work, may capture easily accessible retail and entertainment or the area that determines public good quality. Other difficulties in delineating neighborhoods include our limited knowledge of how housing searchers subdivide a city and the formidable investment required to create a granular geography with coverage of a large country.

Likely as a result of these difficulties, few neighborhood geographies are available in the United States. The vast majority of research uses the census tract, and options are especially limited at broader resolutions (Lee et al. 2008). This reliance on a single scale may limit our understanding of neighborhoods in several ways. Prior work suggests neighborhood effects may depend on both local and broader contexts (Billings et al. 2022), that important drivers of segregation may operate at a level above the tract (Davis et al. 2025), and that neighborhood policies can spill out of their targeted areas (Hanson and Rohlin 2013; Ziff 2025). A broader geography would enable further work on these topics, as new microdata made hyper-local analyses possible (Logan and Parman 2017; Bayer et al. 2025). In addition, a larger unit could improve the precision of local model estimates and data aggregates, such as home price appreciation (Bogin et al. 2019) or economic opportunity (Chetty et al. 2025), and solve some practical problems in the estimation of quantitative spatial or neighborhood choice models (Dingel and Tintelnot 2026; Bayer et al. 2016).

In this paper, we develop a revealed preference approach to neighborhood delineation and use it create a new geography with an average population of 40,000. We call these units "districts," following Jacobs (1961), and we construct them as discrete combinations of tracts, nested within

counties. A tract-to-district crosswalk and set of interactive maps are available at the authors' website. We build the districts using a network algorithm that groups together tracts that see high rates of cross-migration, which we observe using nearly comprehensive address history data from the Census Bureau. This allows us to quickly create the geography in different metropolitan areas and time periods, in contrast to traditional labor-intensive methods. After showing that districts align well with popular perception, physical barriers, administrative boundaries, and demographic discontinuities, we illustrate their usefulness for studying neighborhood sorting and aggregating spatial data.

We begin the paper with a review of qualitative work on neighborhood scale, prior attempts to create new neighborhood geographies, and the existing substate geographies. We then turn to a simple model that illustrates why revealed preference might work in this setting. In the model, individuals face a moving cost to changing districts and have persistent tastes for neighborhood characteristics. These factors generate the key prediction that underlies our empirical approach—people should disproportionately move within districts. Basic statistics using the address history data support this model implication. Migration from an origin tract is highly concentrated in the tract's most frequent destinations, which also tend to be its closest neighbors. However, even among the nearest neighbors, there is a left tail that receives few migrants from the index tract, consistent with discrete neighborhood boundaries.

This feature of migration within cities, which has not been documented previously, both supports the existence of popularly known broad neighborhood areas and opens the door to backing out their boundaries from observed location choices. This sort of revealed preference approach has been used to measure difficult-to-observe quantities such as university quality (Avery et al. 2013) and job amenities (Sorkin 2018). It has significant benefits—it is scalable and does not impose much structure on what a neighborhood should look like—but it is also computationally difficult in the neighborhood delineation setting. We must corral hundreds of millions of moves and select a partition of each city's tracts from a nearly infinite set of possibilities.

To efficiently implement the idea, we turn to community detection algorithms, which group nodes of a network into strongly interconnected modules. The Infomap algorithm, developed in Rosvall and Bergstrom (2008), is particularly well-suited to migration applications because it accommodates weighted, bidirectional connections. To apply Infomap, we first convert the migration flows within

each county into a network in which nodes are census tracts and the connection between Tract A and Tract B is a function of the number of migrants from A to B between 2010 and 2019. We then use the algorithm to partition each network into modules with high rates of cross migration, which form our districts. Perhaps the most important specification decision is the target district size: we aim for a mean of eight to 10 tracts and a range of five to 20 tracts. This roughly splits the difference between counties and tracts (adding the maximum additional spatial information), corresponds with the observed concentration of migration flows, and aligns with some popular colloquial neighborhood definitions (such as Chicago's Community Areas). We also require that districts be contiguous and have population over 12,000.

We estimate districts in the 100 largest Core-Based Statistical Areas (CBSAs). The resulting tract-to-district crosswalk is available at https://evanemast.github.io/district-maps/, along with interactive maps that overlay districts on base maps of streets and geographic features. The interquartile range of district size is six to 12 tracts, corresponding to 27,000 to 54,000 people. Even though the algorithm only uses spatial information in rare corner cases, districts are generally compact and regular in shape. Turning to internal validity, the average migration connection between a pair of tracts in the same district is much stronger than between an equidistant pair of tracts not in the same district. In addition, larger districts have less concentrated migration flows, suggesting size variation is at least partially data-driven.

District boundaries have a striking coincidence with important breakpoints in cities. Consistent with a growing literature on highways and cities (Bagagli 2023; Brinkman and Lin 2024; Wang 2025), we find that adjacent tracts separated by an interstate are disproportionately likely to be placed in different districts, especially in dense areas. Similar patterns emerge for rivers and railroad tracks. District boundaries also tend to coincide with municipal and school district borders, consistent with sharp breakpoints in public good quality (Schönholzer 2023), as well as with abrupt demographic shifts. Since many district borders are formed by these salient features, they tend to match recognizable and cohesive areas, as shown in Figure 1. In cities, districts often resemble popularly known neighborhoods, as the figure shows for McKinley Park in Chicago and the West Village in New York City. In suburbs, they often coincide nearly exactly with towns or school districts, as the figure shows for Dedham, MA and Lower Merion Township, PA. These examples also give a sense of districts' scale and the spatial attributes they may capture or endogenously

determine. Districts seem to be approximately the scale that people have in mind when answering the question, "Where do you live?"

Overall, we argue that districts provide a level of quality comparable to what would be produced by costly, labor-intensive convenings of local experts. Revealed preference identifies the important breakpoints in cities, out of the huge set of features that could potentially form neighborhood boundaries, by selecting those that measurably influence migration. Of course, like census tracts or any comprehensive neighborhood geography, districts are not perfect. We discuss common sources of irregularities, such as weakly defined neighborhoods or odd census tract boundaries, and argue that they are unlikely to be important in empirical applications.

The most concrete contribution of this paper is the construction and public release of this new neighborhood geography, which fills the large gap between tracts and counties and has many potential applications. The underlying conceptual innovation is a new method for applying the economic principle of revealed preference to neighborhood delineation. Since the method relies on an off-the-shelf algorithm with low computational burden, it can easily be replicated in other settings with administrative address history data. To our knowledge, no prior work has constructed a national geography or identified neighborhoods using revealed preference, although others have drawn neighborhoods in small sets of cities (Grannis 1998, 2005; Clapp and Wang 2006; Sampson and Sharkey 2008; Spielman and Logan 2013). Recent work has also used algorithmic methods to delineate space for other purposes, notably contemporaneous work by Ziff (2025), who uses migration flows to estimate housing submarkets in Chicago.[1] More generally, researchers have produced a growing number of large-scale data sets on urban topics including zoning stringency (Bartik et al. 2025) and amenity access (Cook 2025).

The final sections of the paper describe two illustrative applications. First, districts can be used to incorporate broader neighborhood areas in studies of sorting and segregation, which have largely relied on census tracts (Lee et al. 2008). We show that tract-level analyses miss important elements of racial segregation that occur at a broader scale: over 87 percent of variation in tract racial composition is across districts, and segregation indices are nearly identical at tract and district

[1] Researchers have also defined labor markets based on job transitions (Schmutte 2014; Nimczik 2023; Fogel and Modenesi 2023), constructed regions based on commuting flows (Hamilton and Rae 2020), and identified mobility barriers faced by housing voucher holders (Garboden 2021). Daepp (2022) applies Infomap to migration between towns and hand-constructed neighborhoods to construct urban planning regions in Massachusetts.

levels. While a small literature has used buffer regions and largely found similar results (White 1983; Echenique and Fryer Jr 2007; Reardon et al. 2008), districts respect stark neighborhood divides that buffers do not and allow researchers to avoid an arbitrary decision on buffer radii.

However, there is significant within-district variation in socioeconomic indicators like tract median household income, suggesting that tract characteristics do not completely characterize the local environment. District characteristics are correlated with measures of a tract's quality even conditional on the tract's other characteristics, implying that tract-level analyses may not fully capture the relevant area for neighborhood effects (Chyn and Katz 2021) or racial disparities in neighborhood quality (Logan 2011; Bayer et al. 2021; Aliprantis et al. 2022). For example, the average majority-Black tract is in a district with a poverty rate 7 percentage points higher than a majority-white tract with similar tract-level characteristics. We conclude the section by discussing possible implications for models of neighborhood choice and change.

Last, we show that districts are a useful level of spatial aggregation. Aggregating the Opportunity Atlas estimates of tract intergenerational mobility (Chetty et al. 2025) to the district level shrinks standard errors by a factor of three while introducing minimal aggregation bias. Districts preserve 83 percent of signal variance, and average mobility in a tract's district is actually a better predictor of the tract's mobility estimate in a holdout sample than the tract's own baseline estimate. This suggests that districts could help solve recurring sample size issues in the literature on local opportunity. For example, Chetty and Hendren (2018) must estimate causal "neighborhood" effects at the county level because their empirical strategy is underpowered at the tract level. Districts are large enough to support their methodology, which would yield causal estimates for recognizable neighborhoods. Finally, we show that using district as the unit of choice in microfounded models offers both substantive benefits and practical advantages related to dimensionality and granularity.

1 Background and Data

1.1 Prior Work on Neighborhood Definition and Scale

Galster (2001) provides a useful conceptual definition of a neighborhood: a contiguous area with a similar set of spatially based attributes or externalities. This emphasis on common attributes fits naturally with many economic models in which local characteristics vary at the neighborhood level

(in some cases because they are determined by a neighborhood's residents). The definition also easily accommodates multiple neighborhood scales. Some attributes of a person's neighborhood environment are quite local and vary from one block to the next. Others vary more broadly, perhaps because people travel farther to access them or because they are determined by the people living in a broader watershed. Survey work has found wide variation in self-reported neighborhood size, suggesting that people perceive different neighborhood scales in different contexts.[2]

A very small neighborhood area, such as the census block, may determine local attributes like the neighbors that one runs into regularly, nearby property upkeep, or noise levels. Suttles (1972) describes the smallest neighborhood scale as the area in which children can play unsupervised, while Jacobs (1961) discusses an overlapping set of "street neighborhoods" that may be as small as a single block. Prior work points to important peer effects and strong network connections with these close neighbors (Bayer et al. 2008).

A broader neighborhood area shares a different set of attributes. Jacobs (1961) describes "districts" as collections of potentially heterogeneous street neighborhoods which people frequently travel between for everyday tasks. They often feature popular centers of use such as parks, civic facilities, and commercial areas, and their borders are typically defined by physical barriers (such as a busy street) that prevent one street neighborhood from overlapping with the next. Tach et al. (2016) discuss a number of amenities that are likely located in one's district but not their block: schools, grocery and retail shops, parks, transit stops, non-profit organizations, recreation centers, and libraries. Measuring these amenities at the block level would misstate both what is easily accessible from a given location and the watershed that determines the quality of these features. These larger neighborhood areas are also frequently the scale at which neighborhoods develop recognizable names, community groups, and related social networks.[3]

A researcher may want to work at either a small or large neighborhood scale, depending on which local attributes, externalities, and social or economic processes are relevant to an application.

[2] Talen et al. (2025) collect 5,000 hand-drawn neighborhood boundaries in Chicago and report a population interquartile range of 11,000 to 41,000. Pebley and Sastry (2009) report that 40 percent of their Los Angeles sample consider their neighborhood to be either the area within a 15-minute walk or a larger area, while 35 percent consider it to be the block they live on. (They also find considerable variation between respondents in the same census tract.) Coulton et al. (2013) report that the 75th percentile of their survey responses is about twice the land area of the average census tract, while the 25th percentile is only 20 percent of the tract average size.

[3] For example, Luttmer (2005) documents a negative relationship between an individual's happiness and their relative income using neighbors within the same Public Use Microdata Area as the reference group.

However, the set of purpose-built neighborhood geographies is limited, especially if one desires national coverage. Several recent projects construct neighborhoods for particular cities based on large-scale surveys (Woodruff 2017; Buchanan 2023; L.A. Times 2025; Talen et al. 2025). This approach results in neighborhoods on the district scale, with mean neighborhood populations across these projects ranging from 23,000 to 40,000. Other researchers have developed methodologies based on first principles. Grannis (1998) and Grannis (2005) delineate neighborhoods as areas that are connected by small neighborhood streets, and Clapp and Wang (2006) divide West Hartford, Connecticut into neighborhoods that minimize variance in property values. Some researchers, including Spielman and Logan (2013) and Sampson and Sharkey (2008), have clustered together individuals or census tracts with similar demographic characteristics. However, none of these efforts has been completed at the national level, likely due to the prohibitive cost of scaling methodology (particularly surveys) to cover the entire country.

1.2 Existing Substate Geographies

While researchers have not created a national neighborhood geography, the federal government produces a number of substate geographies that are often used as neighborhood units. Table 1 describes the important features of counties, Public Use Microdata Areas (PUMAs), zip code tabulation areas (ZCTAs), census tracts, census block groups, and census blocks.

Tracts, block groups, and blocks all capture much smaller areas than districts. These geographies are cleanly nested in each decennial year, although the Census Bureau does change them over time to accommodate population change. The tract, which contains approximately 4,000 people, is the primary scale used in neighborhood research. It is popular in part because researchers have developed crosswalks that harmonize boundaries from 1970 to the present (Logan et al. 2014), and in part because the Census Bureau releases much more demographic information for tracts than blocks or block groups. In addition, because tracts are drawn to be easily identifiable in the field, their borders often coincide with permanent physical features like highways, waterways, and railroad tracks. However, there is no reason to think that tracts represent a particularly important neighborhood scale. They were created gradually over the first half of the 20th century to provide statistical information on subareas of a city, with an initial emphasis on disease measurement. To our knowledge, the only stated rationale for their size is a minimum level of statistical precision. However,

it is noticeably similar to the “neighborhood unit" concept, which was developed contemporaneously by urban planners and targeted a population of 7,000, justified as the minimum scale to sustain an elementary school (Perry 1929).[4]

The existing options for delineating a neighborhood larger than the tract have significant limitations. In Section 4, we show that these geographies frequently cross important physical barriers and sharp demographic discontinuities, potentially creating significant measurement error. First, ZCTAs approximate postal zip codes, which were drawn to optimize mail delivery and were never intended to reflect neighborhood boundaries. They vary widely in size and overlap unevenly with tracts, with over a quarter contained entirely within a single tract. (For example, Chicago's enormous Merchandise Mart office building famously had its own zip code until 2008.) In contrast, PUMAs are drawn with demographic similarity in mind, and they align with census tract and county boundaries. However, their minimum size is 100,000 people, and over 25 percent contain more than 240,000, making them significantly larger than the neighborhood sizes reported in surveys. In addition, PUMAs can be spatially discontiguous and are drawn by state-level agencies that use heterogeneous methodologies (Siordia and Wunneburger 2013).

Finally, researchers will sometimes use a circular buffer region around an individual's residence in place of neighborhood geography. This defines a set of overlapping neighborhood environments, rather than partitioning space into unique units. Buffers offer some benefits—for example, they may capture the environment of a person living near the edge of a spatial unit better than the unit's boundaries. However, the approach does not account for sharp breaks between neighborhoods, which could lead to buffer regions that include areas that the index person would not consider their neighborhood, and researchers must arbitrarily choose a radius length.

1.3 Data

Our primary data source is the Census Bureau's Environmental Impacts Frame (EIF), which provides annual, address-level residential histories for nearly all U.S. residents from the late 1990s onward. Addresses are drawn from an array of administrative sources, with a priority ordering based on

[4] Jacobs (1961) strongly criticized this neighborhood scale: “The ideal neighborhood of planning and zoning theory, too large in scale to possess any competence or meaning as a street neighborhood, is at the same time too small to operate as a district. It is unfit for anything. It will not serve even as a point of departure. Like the belief in medical bloodletting, it was a wrong turn in the search for understanding.”

each source's accuracy and reliability. The most frequently used source is annual 1040 and 1099 tax filings, but addresses can also be taken from the Medicare Enrollment Database, HUD records on housing assistance programs, and the USPS National Change of Address file. Voorheis et al. (2023) describe the construction of the EIF in detail.

We use the EIF to construct a data set consisting of the number of moves from each possible origin tract to each possible destination tract from 2010 to 2019.[5] We define a move from Tract A to Tract B as a person who is observed in Tract A in year t and in Tract B in year $t+1$. In the estimation of the publicly released version of the districts, a small amount of noise is added to the move count variable to protect confidentiality. We show later that this does not meaningfully affect results. Finally, we use the 2010 vintage of census tracts and draw demographic information from the Longitudinal Tract Data Base (LTDB) developed in Logan et al. (2014).[6]

2 Methodology

2.1 A Revealed Preference Approach to Delineating Neighborhoods

Why might migration patterns identify neighborhood boundaries? Suppose that a neighborhood construct like districts exists, that each district has attributes that affect the experience of living in it, and that the district map is known by individuals living in a city. Consider an individual who currently lives in tract c_{-1}, and let each district d be composed of a collection of tracts. The individual chooses a tract for the next period t to maximize the indirect utility function:

$$u_{ict} = \alpha_i x_c + \beta_i x_{d(c)} + \gamma \mathbb{1}(d(c) \neq d(c_{-1})) + \epsilon_{ict} \tag{1}$$

where x_c is a vector of tract characteristics, $x_{d(c)}$ are the characteristics of the tract's district, and α_i and β_i are individual-specific, time-invariant preferences over those characteristics. Since $d(c_{-1})$ is the individual's baseline district, γ is a moving cost from changing districts.

All else equal, the individual is clearly relatively likely to choose a tract in the same district as their initial tract. This first avoids the moving cost γ. In addition, since the individual chose $d(c_{-1})$

[5] Note that because migration rates vary widely with demographics, some subgroups will account for a disproportionate share of moves. Appendix Table A.1 shows that renters account for 64 percent of moves, and individuals aged 19-29 account for 26 percent. In contrast, people over age 65 account for less than 5 percent.

[6] Section A.3 of the Appendix provides more details on data sources and sample definitions.

in an earlier period and preferences are stable, the characteristics of this district likely generate high utility. Finally, if tracts within a district are relatively homogeneous, a similar argument applies to tract characteristics. Thus, migration flows should be disproportionately clustered within districts. This implies that if a district construct exists and influences location decisions, an econometrician could back out its boundaries by identifying clusters of tracts with high cross-migration.

This stylized model provides a foundation for our revealed preference approach. However, it relies heavily on the assumption of time-invariant preferences over neighborhoods. This may be the case for some moves. Perhaps a renter's lease has expired, or a person wants to change housing unit size or quality. On the other hand, many people move precisely because their taste for neighborhoods changes, perhaps due to changes in their income, household composition, or age.[7] Fortunately, our methodology only requires that an excess mass of within-district moves can be detected empirically, not that all moves be within a district. This condition holds in practice; in part because moves driven by changing neighborhood preferences are relatively dispersed (since preferences can change in many directions) and thus do not create many large flows.

There are a number of benefits to a revealed preference approach. First, unlike methods that cluster together demographically similar areas, it does not impose any assumptions on racial or economic homogeneity within neighborhoods. This is especially important because neighborhood geographies are often used to measure sorting and segregation. Our method requires only that neighborhoods have some shared attributes that influence individual preferences. Second, it can be implemented using administrative data on migration, which is available in many settings. In contrast, survey-based approaches require costly and specialized data collection, and they often struggle with low responses rates in poor and heavily minority areas (Talen et al. 2025).

2.2 Empirical Patterns in Migration Flows

Before attempting to back out district boundaries, we use the EIF to show that raw migration flows across census tracts are consistent with the existence of a broader neighborhood concept that influences location decisions. Figure 2 shows the share of within-county migrants from the average

[7] Responses to the Current Population Survey suggest that migration is most frequently driven by a desire to change housing units or a change in household composition, motivations that could conceivably lead someone to search within their district or in a new district. Only about 10 percent of moves are driven by factors that almost certainly make a person want to change districts: seeking a better neighborhood and reducing commute.

tract to each of its 50 strongest migratory connections and its 50 closest tract neighbors by centroid distance. Migration is strongly concentrated along both dimensions. The strongest connection receives over 6 percent of departures from the index tract, and the closest neighbor receives more than 4 percent. These shares rapidly decline, with the 25th ranked tract on either metric receiving less than 1 percent. Moreover, the curves are convex, with departures shares sharply decreasing until about the tenth-ranked tract, when the slope begins to flatten.

Migration is not simply concentrated among the closest tracts; Panel A of Appendix Table A.2 illustrates that some tracts exchange relatively few migrants even with tracts they directly border. Slightly over 7 percent of adjacent tracts are not among the index tract's 20 strongest connections. This is consistent with neighborhood boundaries that affect migration, and we show after estimation that district borders are drawn precisely at these places where migration connections are weak. Panel B instead shows the distance rank of close migratory connections, with 80 percent of the index tract's 10 strongest connections coming from its 20 closest neighbors.

2.3 The Infomap Algorithm

Inverting migration flows to district boundaries at scale is not straightforward. The number of possible partitions of the tracts in a large metropolitan area is nearly infinite, and the EIF data contains hundreds of millions of moves.

One approach would be to estimate Equation 1, searching for the parameters and district partition that maximize the likelihood of observed tract choices. While this is theoretically attractive, there are a number of practical difficulties. Foremost is the huge number of possible partitions. In addition, the model contains many other parameters that are difficult to estimate and may interact with the district boundary estimates, but are not central to the goal. We instead take an algorithmic approach. This simplifies estimation immensely, resulting in quick run-times that allow for iterative testing and development. It also makes our estimation process much easier to implement in other settings than a state-of-the-art neighborhood choice model.

We use the Infomap algorithm developed in Rosvall and Bergstrom (2008), which is among a large class of community detection algorithms. Infomap takes a weighted and directed network as an input. Roughly speaking, it partitions the nodes of the network into the set of modules that minimize the expected frequency at which a random walker surfing the network will transition

between modules, subject to a penalty for the modules becoming too large.[8] In our application, nodes are tracts, connection strength is determined by migration frequency, and the modules created by the algorithm are the districts. The algorithm's objective function will thus lead it to group together tracts that have relatively high cross-migration. (Variation in cross-migration would also drive the district partition estimates in a maximum-likelihood estimation of Equation 1.)

There are several reasons that Infomap is ideally suited for our application, although other community detection algorithms could likely be used to obtain similar results. First, the revealed preference argument suggests that pairs of tracts within a district should see migration in both directions, and Infomap accommodates a directed network. This minimizes the influence of moves that are due to changing neighborhood preferences, which tend to be unidirectional (e.g., suburbanization after having children). Second, Infomap has a high resolution, enabling it to construct smaller modules than competing methods like stochastic block models (Kawamoto and Rosvall 2015). This is important because our desired district size is small relative to large counties. Third, Infomap has a low computational burden. Running Infomap on Los Angeles County, the most populous in the United States, only takes about a minute on a standard personal computer.

3 Estimation

3.1 Choosing Target District Size

We use three sources of information to select a target size: prior neighborhood definitions, patterns of tract-tract migration connectivity, and diagnostic tests of the estimation output.

The survey-based neighborhood delineation efforts discussed in Section 1.1 yielded mean sizes of 30,000 (Chicago), 33,000 and 44,000 (Los Angeles), 23,500 (New York City), and 32,000 (Boston),

[8] Rosvall et al. (2009) provides an intuitive description of the Infomap algorithm, which is built on concepts from information theory. To summarize, suppose that we started a random walk at a particular node and allowed it to move along the network. The walker draws the next node from the set of nodes connected to its current node, with probability proportional to the connection strength to each node. The goal of the Infomap algorithm is to construct a binary coding system that can describe the nodes visited by the random walk in as few bits as possible (in expectation). It constructs a two-part coding system, in which modules are assigned one set of codes, and nodes within each module are assigned a separate set. There is also a code that indicates when the walker switches between modules, so that the coding system can use repeated codes for the nodes within each module and still uniquely describe each walk. The benefit to the two-part system is that short binary codes (e.g., one bit for the number 1) can be used repeatedly for nodes in different modules, avoiding the use of long codes for larger numbers (e.g., nine bits for 300). The goal of minimizing expected description length pushes the algorithm to construct strongly interconnected modules that a walker is relatively unlikely to exit, reducing the number of module transitions that must be encoded. On the other hand, having more modules allows the algorithm to use more short node codes.

and expert-drawn geographies like Chicago's Community Areas and New York City's Neighborhood Tabulation Areas fall in a similar range. Colloquial neighborhood boundaries in smaller cities tend to be significantly smaller (Hartley and Rose 2023), likely because a broader geography is not useful shorthand when a metro area contains only a handful of district-level areas. However, the local attributes determined at the district level should be similar across cities of different sizes, so we take the 25,000–45,000 range from larger cities as our primary reference point.

The raw migration data provides independent evidence pointing to a similar scale. Figure 2 shows that departures from a tract are concentrated among its five to 20 closest connections. This suggests that groups of tracts in this size range may form distinct neighborhood units containing highly-substitutable subareas (although sharp drop-offs in connectivity at district borders will not appear in this figure due to averaging across tracts).

Together, these observations suggest a rough initial target around 10 tracts. To finalize the size range, we ran the estimation routine (described in full in the next subsection) under several parameterizations and computed two diagnostic tests. When estimated districts are too small, we expect that some boundaries will reflect weak differences in migration connectivity that arise from sampling noise rather than meaningful separation. We assess this by estimating districts on repeated 10 percent subsamples of the EIF and measuring the stability of the resulting boundaries. When estimated districts are too big, they may contain subareas that are not strongly connected by migration and should instead form separate neighborhoods. We assess this by calculating the share of estimated districts that have pairs of tracts that cannot be linked by a chain of common within-district moves.

The parameterization that performed best on these diagnostics yields a mean district size of nine to 10 tracts and a distribution contained within five to 16 tracts. In addition to the rationale above, this size range sits at a convenient midway point between counties and tracts. The median county in our sample contains 25 census tracts, and the 75th percentile has 95 tracts, so this target size breaks the typical county into about five to 15 districts that each contain five to 15 tracts. This is appealing because if a district were only slightly smaller than a county, or only slightly larger than a tract, it would add little information about the local environment.

3.2 Converting Data to Network Form and Estimating Infomap

In order to apply Infomap, we convert the tract-to-tract transition data (created from the 2010 to 2019 EIF files) into a weighted and directed network in which tracts are nodes and the connection strength is a function of the number of migrants. We construct a separate network for each county that includes only tract pairs within it (implying that only within-county moves, about 60 percent of all migration, enter the estimation), and estimate Infomap independently on each network. This ensures that counties nest districts and reduces computational burden in large CBSAs.[9] We do not include tracts with a population below 500, which see little migration and are generally non-residential areas like airports, parks, or industrial agglomerations.

The primary degree of freedom in this transformation is the definition of connection strength between a tract pair. Different definitions yield significantly different results, especially for district size. When strength is very concentrated in the tightest connections (e.g., number of migrants squared), Infomap returns small districts, as a random walk is unlikely to exit even very small modules. In contrast, when the connection strength is more diffuse (e.g., square root of number of migrants), districts are large, as the random walker will frequently exit even large modules.

To choose a connection strength function, we estimate a number of specifications and consider their performance on several metrics. First, we consider basic measures like the share of districts within the target range of five to 20 tracts or that are initially discontiguous. Second, we assess overfitting by estimating districts using 10 percent subsamples of the EIF and calculating the stability of boundaries. Finally, we estimate the share of districts that contain subareas that are not strongly connected by migration, which suggests that distinct districts are mistakenly being combined. We deliberately do not consider demographics in this informal specification search.

The results from each candidate specification are described in detail in Section A.1 of the Appendix. The crucial insight is that basing connection strength on a destination tract's rank among an origin tract's possible destinations, rather than raw counts of migrants, standardizes the strength distribution and allows the algorithm to perform well across tracts with different populations and demographic characteristics. Specifically, we divide the number of migrants from tract A to tract B by the total number of migrants to B, producing a measure of B's popularity that accounts for its

[9] We set all counties with a population below 30,000 as a single district to avoid some odd corner cases.

size. We then rank all of A's destinations according to this variable, with one being the highest. Finally, denoting this rank as R_{ab}, we define connection strength as $\frac{1}{(R_{ab}+2)^X}$ if $R_{ab} \leq C$ and 0 otherwise.[10] This functional form performs well on all metrics, and the exponent X and connection cutoff C allow for finetuning, with higher exponents and tighter cutoffs leading to smaller districts. We set $X = 2$, $C = 10$ and show later that tweaks have little effect.

With the connection weights defined, the network representation of the tract-tract flow data is complete. We apply the Infomap algorithm to each county's network and obtain a crosswalk mapping each tract to a rough set of districts. To arrive at the final set of districts, we regularize this initial output in two ways. First, we split districts that are not spatially contiguous into their contiguous subcomponents.[11] Second, when a district has a population below 12,000, we combine it with the adjacent district to which it has the strongest migratory connection. About 8 and 2 percent of tracts are in initially discontiguous or low-population districts, respectively.

4 District Boundaries and Validation

4.1 Overview

In this section, we illustrate that the district geography captures distinct migration patterns and aligns well with physical, administrative, and demographic discontinuities. Revealed preference identifies the important breakpoints in cities by selecting features that affect migration. While the best way to assess district quality is to examine the interactive map of a familiar city, available at https://evanemast.github.io/district-maps/, we include several examples here. Figures 3 and 4 show districts in Chicago city proper and the surrounding CBSA, respectively. Appendix Figures A.1 to A.3 show Manhattan, Washington, D.C., and the Houston metropolitan area.

The maps exhibit some encouraging visual patterns. Districts are generally spatially regular and close to convex—they are not shaped like a noodle or an octopus. Consistent with patterns in density and travel speed, districts gradually increase in size as distance to the city center increases. While it is difficult to quantify, districts also typically pass the eye test. Figure 1 shows four illustrative examples. McKinley Park in Chicago has borders that are clearly defined by a combination of

[10] We shift the rank by two in order to flatten the gradient among the strongest connections.
[11] We allow "queen" or "kitty corner" adjacencies, in which tracts only touch at a single point in space.

railroads, industrial yards, and waterways—district borders capture this neighborhood exactly. The district corresponding to West Village in New York City closely matches survey work from Buchanan (2023). Last, the suburban municipalities of Dedham, MA and Lower Merion Township, PA are each exactly one district. Somewhat more systematically, Appendix Figures A.4 and A.5 show that districts align well with Chicago's Community Areas and New York City's Neighborhood Tabulation Areas in many cases, particularly at important physical barriers.

Districts also tend to fall around the target size range, as shown in Table 2. The average district contains about nine tracts and 40,000 people, and the interquartile ranges for these variables are six to 12 and 27,000 to 54,000. The lower panels describe district sizes in towns with top quartile population density versus others. Population and tract counts are similar in the two groups, but districts in high density places have a much smaller land area.

We find strong evidence that district boundaries are driven by migration patterns—Infomap successfully implements the revealed preference idea. A tract is placed in the same district as its strongest connection 95 percent of the time, versus only 14 percent for its 20th strongest (Appendix Figure A.6). In addition, there are large connectivity differences between equidistant tract pairs that are in the same district versus those that are not. Figure 5 shows that when the nearest tract is in the same district, it is the index tract's fourth strongest connection on average. When it is in a different district, it is 30th. Last, variation in district size is at least in part driven by the data. Appendix Figure A.7 plots the share of departures to a tract's strongest connections separately for districts with fewer than seven versus more than 14 tracts. Nearly 8 percent of departures in small districts are to a tract's strongest connection, and 26 percent are to the five strongest. In larger districts, the corresponding figures are only five and 18 percent.

Boundary estimates are also relatively stable. We estimate districts separately on 10 percent subsamples of the EIF (drawn without replacement) and calculate, for each pair of adjacent tracts, the number of subsamples in which the two tracts are placed in different districts. Appendix Figure A.8 shows that half of tract pairs are never a district border, while 14 percent are a border in every subsample. Only 15 percent of tract pairs receive a border in between three and seven subsamples (inclusive). In addition, the noise infusion required to disclose the publicly available tract-to-district crosswalk has little effect. About 94 percent of adjacent tract pairs have the same district-border status in the baseline and the publicly available version, including 99 percent of "hard" tract borders

that appear in all 10 subsamples. Finally, we also created districts based on migration during from 2000 to 2009. About 90 percent of adjacent tract pairs maintain the same district relationship in this version and the primary estimation. Section A.2 of the Appendix discusses stability across small changes to the specification.

4.2 Alignment with Physical Barriers and Local Government Borders

As an initial illustration of how physical barriers tend to form district borders, Figure 6 overlays interstate highways and rail lines on a map of districts in Chicago. Example boundaries include Interstate 55, which extends southwest from the central business district and the parallel rail line a mile to the north of I-55.[12] In Table 3, we generalize this example by considering all adjacent tract pairs and all limited-access highways. In pairs separated by a highway, the two tracts are in different districts 48 percent of the time, versus a rate of only 30 percent among other pairs. Highways are even more likely to form district boundaries in dense areas, where they are a larger barrier to common modes of transport (Bagagli 2023; Brinkman and Lin 2024). Among tract pairs with above-median population density, 63 percent of those separated by a highway span a district border, versus only 36 percent of pairs that are not separated by a highway.

Municipal borders are also important determinants of boundaries, particularly in suburban areas where towns may be near the target district size. For example, in the Boston suburbs, the towns of Brookline, Revere, and Needham are each exactly a district. This likely results from preferences over public goods, particularly families who want different housing but do not want to change school districts. The second panel of Table 3 shows that 56 percent of adjacent tract pairs that span a town border are in different districts, compared to only 25 percent of pairs that are within the same town.[13] Unsurprisingly, much of the importance of town borders appears to be driven by school districts. Sixty-eight percent of tract pairs that span a school district border are classified into different district neighborhoods, versus only 28 percent of pairs that are within the same school district. However, even pairs of tracts that span only one type of boundary—town but not school district, or vice versa—are placed in different districts at an elevated rate, suggesting that both schools and public goods generally influence migration.

[12] In contrast, Appendix Figure A.9 shows that ZCTAs in Chicago do not align with these features.
[13] We do not include pairs that span a county border, since we impose that these locations are district borders.

4.3 Alignment with Demographic Differences

The desired alignment between districts and demographic changes is nuanced. We would usually like to see a district border where there is a large change in racial or economic composition, since the demographic change likely indicates a break between neighborhoods. However, some neighborhoods may span these divides, and a benefit of the revealed preference approach is that it does not impose boundaries at these locations. We would also like to see similar adjacent tracts be placed in the same district, but neighboring districts could certainly be demographically similar and yet distinct.

To illustrate these ideas, the left panel of Figure 7 overlays district borders on a map of tract racial majority in Chicago (taken from the 2010 census). The coincidence of borders and racial composition is striking, particularly along the boundaries between the majority-Black west side and the majority-Hispanic northwest and southwest sides. Districts are generally composed of a similar set of tracts, sometimes with some bleeding over from an adjacent district that may represent ongoing neighborhood change or racial compositions that are close to a 50 percent cutoff.

In contrast, alignment with demographics appears to be a shortcoming of ZCTAs, as shown in the right panel of Figure 7. Many more units contain areas with very different racial composition, with perhaps the most striking example being the 60623 zip code, which lies along the western border of the city just south of its midpoint. This zip code combines majority-Black North Lawndale and majority-Hispanic Little Village, which are in reality sharply divided by a rail viaduct running through the middle of the area. ZCTAs can also split coherent neighborhoods, such as Hyde Park around the University of Chicago (the majority-white and no-majority area to the southeast).

Moving beyond examples, Table 4 shows that districts are less likely to contain a large demographic discontinuity than ZCTAs or buffer regions. We define the buffer region for each tract as the set of tracts that border it, as in Logan (2011). We define a large discontinuity as a pair of adjacent tracts that are in the top 5 percent in their difference in white share or log median income. Because the likelihood of a discontinuity mechanically increases with size, we stratify by number of tracts contained in the unit and do not include units larger than 12, as few ZCTAs or buffers exceed this threshold. At the median district size of 9-10 tracts, 13.8 percent of districts contain a discontinuity in white share, versus 26.2 percent of ZCTAs and 27 percent of buffers. Relative differences are

similar across sizes, and districts also perform better on income, although by a smaller margin.[14] Appendix Figure A.10 shows a complementary fact: districts are significantly less likely than ZCTAs to place a border between a pair of adjacent tracts with a similar white share, as well as slightly more likely to place a border between a pair with very different white shares.

4.4 Weaknesses and Imperfections

The transparency of district boundaries makes their imperfections more visible than most estimates from economic models. In this subsection, we describe common sources of irregularities or errors. On the whole, we believe that they are unlikely to meaningfully undermine the usefulness of districts and are likely unavoidable in any attempt to define neighborhoods at scale.

First, because districts are built from census tracts, they inherit some of their limitations. Tracts can take highly irregular shapes, especially when they include unpopulated areas like industrial zones or parks. These spaces are often not strongly associated with any neighborhood and may in reality serve as "thick" boundaries. In addition, while tracts generally align well with major physical features, there are exceptions. A tract may span a highway, or a major street may serve as a boundary between a series of tracts, then fall within a single tract, then resume its role as a dividing line. Some tracts are drawn to conform with irregular municipal boundaries. Appendix Figure A.11 shows examples of districts influenced by each of these cases. Using the tract as the building block also prevents us from separating distinct subareas within a tract. However, we believe that the convenience of a unique mapping from tracts to districts outweighs these costs.

A second class of imperfections arises when neighborhood boundaries are not well defined. For example, consider two demographically similar districts that are separated by a large street but see significant cross-migration. The algorithm may assign a tract on one side of the street to the other side's district, producing a zigzag boundary rather than a straight line. In general, these cases are unlikely to affect estimates of district composition, as they occur precisely where small errors in the boundary are unimportant.

Third, we do not assign the 1 percent of tracts with fewer than 500 people to a district. While very few people live in these areas, they can be spatially large or prominent (e.g., the office district

[14] We do not include PUMAs in this table because they are much larger, with over 90 percent containing more than 20 tracts. Fifty-eight and 64 percent have a white share and median income discontinuity, respectively.

in Midtown Manhattan, the National Mall in Washington D.C., Chicago's O'Hare airport). This may need to be accounted for in applications in which the exact location of jobs is important.

5 Application: Sorting and Neighborhood Choice

5.1 Measuring Sorting and Neighborhood Inequality

In most large cities, it is obvious that racial and economic sorting occur at a broader level than the census tract. However, the vast majority of research on the topic uses the tract as the neighborhood unit (Lee et al. 2008), with a few exceptions examining broader resolutions using buffer regions (e.g., White 1983, Echenique and Fryer Jr 2007, Lee et al. 2008). While buffers are convenient and flexible, there is no clear guidance on their size, and they may cross sharp racial divides or administrative boundaries that determine public good or school quality.[15]

Districts provide an improved way to measure the composition of broader neighborhoods, potentially capturing sorting patterns that have been obvious yet difficult to quantify. Unlike buffers, their size is justified by both empirical migration patterns and prior survey work, and their boundaries align with breakpoints in the local environment. District-level statistics thus have a relatively tight connection with an important set of public goods and amenities, as well as with resident perceptions of their neighborhood.[16] To illustrate potential applications, we use districts to compute racial segregation indices and measure racial disparities in neighborhood quality.

The district-level isolation index for non-white individuals in the average sample CBSA is only very slightly lower than the tract-level index: 0.299 versus 0.308. For Black individuals, the district and tract indices are 0.283 and 0.321. Put differently, the average minority individual lives in a tract that is 64 percent minority and a district that is 62 percent minority, and the average Black individual lives in a 47 percent Black tract and a 43 percent Black district. This implies that tract-level analyses miss important elements of segregation. Nearly identical patterns emerge at a scale 10 times larger, which has implications for segregation's relationship with amenities, public goods, and cross-race exposure, as well as for identifying the drivers of segregation.[17]

[15] Monarrez and Schönholzer (2023) instead show that a large share of racial segregation across neighborhoods can be explained by sorting across local government boundaries.

[16] Buffers offer other benefits, such as capturing differences between the center versus the edge of a district.

[17] While traditional isolation or dissimilarity indices do not accommodate buffer regions, some simple statistics suggest that buffers may understate segregation. Defining a buffer as the set of directly bordering tracts, as in (Logan

The more general relationship between tract and district characteristics is also informative. Table 5 shows the R^2 values from several regressions of tract characteristics on a vector of district fixed effects. Nearly all of the variation in tract racial composition occurs across, not within, districts—R^2 exceeds 0.87 for share white, share Black, and share Hispanic. However, the rightmost columns show that tracts are more likely to share a district with an economically dissimilar neighbor—the R^2 values are 0.696 for log tract income and 0.648 for the poverty rate. To illustrate this pattern, consider majority-Black and high-poverty tracts (setting the poverty threshold at 33 percent to equalize the number of the two types). The mean majority-Black tract is in a district in which 79 percent of the tracts are majority-Black, while the average high-poverty tract is in a district in which only 53 percent of the tracts are high-poverty.

It is somewhat remarkable that districts consist of such racially similar tracts, potentially raising concerns that our method is mechanically producing this feature. This would be true if, for example, all tracts were either all Black or all white, but this is not the case. In fact, our results actually imply that mixed tracts tend to be in similarly mixed districts. For example, Figure 7 shows that the district corresponding to the diverse Rogers Park neighborhood in the northeast corner of Chicago consists almost entirely of tracts with no racial majority. Moreover, the validation in the previous section also suggests that district boundaries match popular perception and important spatial barriers—they do not simply group together tracts with similar racial composition.

Next, we consider racial disparities in the neighborhood environment by regressing district-level economic characteristics on a set of tract-level economic controls, CBSA fixed effects, and dummy variables for tract racial majority. Table 6 shows that majority-Black tracts are located in districts with 6.7 percentage points higher poverty than economically similar majority-white tracts in the same CBSA. This is a large difference, as the mean district poverty rate is 15 percent. Similarly, district income and college share are 20 percent and 4 percentage points lower in the same comparison, suggesting that tract-level estimates understate true racial disparities in neighborhood quality (Reardon et al. 2015; Bayer et al. 2021; Aliprantis et al. 2022).

The importance of this broader resolution of segregation and neighborhood inequality depends on how much the district environment affects individual outcomes and quality of life. Estimating this

2011), the average majority-Black tract is in a buffer that is 64.2 percent Black, versus a 70.5 percent Black district. Over 25 percent of majority-Black tracts are in a district that is more than 11.5 percentage points more Black than its buffer. This is consistent with some buffer regions crossing sharp racial divides.

causal relationship is beyond the scope of this paper, but Appendix Table A.3 shows some suggestive correlations between district composition and tract home values and intergenerational mobility. Tracts that are surrounded by much poorer districts (defined as the top quartile of difference) have 2.4 percent lower home values than tracts in the same county with similar characteristics, and they also have substantially lower mobility rates. In addition, adding district fixed effects explains about a third of the variation that is unaccounted for in a "kitchen sink" specification that uses only tract-level information.

5.2 Models of Neighborhood Sorting

This broader resolution of sorting suggests that models of neighborhood choice and change (i.e., studies of the causes of segregation) could benefit from incorporating mechanisms that operate at a higher aggregation than the tract. Much of the literature follows one of the two models developed in Schelling (1971). First, in the "bounded neighborhood model," agents care only about the composition of a well-defined neighborhood (e.g., Card et al. 2008). Districts provide a way to apply this approach to larger units, which may be helpful in the presence of spillovers across tracts. For example, evaluations of neighborhood revitalization policies have found effects on tracts that border the targeted area (Hanson and Rohlin 2013; Ziff 2025). Second, in the "spatial proximity model," each agent cares about the households within a given radius of their dwelling. This allows for more interesting spatial dynamics, and recent work has taken related approaches to shed light on the spread of majority-Black neighborhoods during the Great Migration (Davis et al. 2025) and gentrification (Guerrieri et al. 2013). Districts could be used to incorporate barriers that limit the influence a tract has on its neighbor. For example, perhaps a person cares about neighbors within half a mile, but not when they are on the other side of a river or in a different town.

6 Application: Spatial Aggregation

6.1 Aggregating the Opportunity Atlas to the District Level

The district geography can be used to improve the precision of numerous tract-level estimates, ranging from simple survey aggregates like ACS estimates of neighborhood median income to model-based estimates like repeat-sale measures of housing price appreciation (Bogin et al. 2019; Contat and

Larson 2024) or housing supply elasticities (Baum-Snow and Han 2024). To illustrate these benefits concretely, we aggregate the Opportunity Atlas (OA) tract-level measures of intergenerational mobility to the district level (Chetty et al. 2025). These estimates are publicly available and widely influential, and follow-up work developed programs that aimed to help housing voucher holders move to high mobility areas in King County, Washington (Bergman et al. 2024). However, the estimates also have large standard errors that make statistical inference difficult, particularly when making ranking-style comparisons among many tracts. For example, Mogstad et al. (2024) show that 93 percent of tracts in King County include both a top-five-percent and bottom-five-percent rank in the 95-percent confidence set on their ranking.

We aggregate the tract-level mobility estimates and their standard errors to the district level using simple population-weighted averages.[18] We focus on estimates of the average adult income percentile rank of children born to parents at the 25th income percentile. A number of results are summarized in Table 7. As expected, district aggregation greatly increases precision—the average standard error on the district estimate is 0.67 percentile points, versus 2.1 for the average tract estimate. Perhaps more surprisingly, the aggregation introduces remarkably little bias. District fixed effects explain 83 percent of signal variance in the tract mobility metrics, implying that district aggregates will rarely misrepresent conditions in their subareas.

Moreover, the remaining within-district variance appears to be mostly noise. We show this using a hold-out sample: tract mobility estimates calculated using the 1985-1989 birth cohorts (versus the 1980-1984 cohorts used in the baseline OA estimates). We predict each tract's mobility in this alternative sample as (1) the tract's baseline mobility estimate and (2) the average baseline mobility estimate in the tract's district. The Mean Absolute Prediction Error (MAPE) in the hold-out sample is actually lower when using the district average: 7.2 percent versus 7.8 from using the tract's baseline estimate. Unsurprisingly, districts' advantage is larger for estimates based on smaller samples. Among Black and Hispanic children, the baseline district estimate generates MAPEs of 10.9 and 10.8 percent, respectively, while the baseline tract estimate produces a MAPE of 14.8 percent for both subgroups. This implies that true within-district variation in mobility is small

[18]The original estimates are based on a parametric model, so this aggregation is an approximation. We test its validity by aggregating tract estimates to the county level using our method and comparing them to the county estimates in Chetty et al. (2025). The correlation between estimates is 0.973, and the correlation between standard error estimates is 0.981.

relative to the reduction in sampling error achieved by district aggregates.[19]

On the whole, district aggregations of the Opportunity Atlas improve precision at little cost of accuracy, generating two immediate benefits. First, districts allow researchers to make precise statements about specific, recognizable neighborhood areas without having to construct the sort of bespoke local geographies that original studies often relied on.[20] We illustrate this by plotting 95 percent confidence sets around the mobility ranking of each of the 36 districts in King County following Mogstad et al. (2024) in Appendix Figure A.12. In contrast to the tract level, no district confidence set spans both the top and bottom five percent of rankings, and none of the bottom tercile of districts includes a top-ten ranking. Second, regression analyses of spatial variation or correlates of mobility using tract-level OA estimates require adjustments for large sampling error; at the district level, sampling error is small enough to be safely ignored.

More substantively, the district aggregation changes our understanding of how intergenerational varies within cities. A map of the tract-level OA estimates often looks like a checkerboard, particularly for subgroup estimates. This suggests that nearly every tract is only a short distance from significantly different levels of opportunity, and that the determinants of mobility are very local and idiosyncratic. However, our results imply that most of this variation is just sampling error, and economic opportunity is actually clustered within recognizable larger areas. This suggests that attributes determined at a broader level—school districts, job and retail accessibility, neighborhood organizations, broad social networks—are likely important.

Finally, districts can solve a recurring problem in both work on local economic opportunity and the neighborhoods literature more broadly: the unit researchers want to study is too small to support the empirical strategies they want to use. An influential example is Chetty and Hendren (2018), who invoke neighborhoods as the object of interest (and use the word in the title), but actually estimate the causal effect of *counties* on intergenerational mobility. Districts provide a set of recognizable neighborhoods that, unlike census tracts, contain enough population to support modern empirical methods. We do not have the data required to replicate Chetty and Hendren (2018), but the county and district population distributions overlap substantially, suggesting that

[19] Aggregating OA estimates to the PUMA level does not yield the same improvements. PUMA fixed effects explain only 66 percent of the signal variance in the tract estimates, and PUMA-level predictions perform worse than tract-level predictions in the hold-out sample.

[20] See the comparison of the Watts and Compton neighborhoods, as defined in L.A. Times (2025), in the introduction of Chetty et al. (2025), or the definition of high-opportunity neighborhoods developed in Bergman et al. (2024).

their empirical strategy could be applied to districts to generate estimates that correspond much more closely to neighborhoods, capture most of the spatial variation in mobility, and have precision comparable to their existing estimates.[21]

6.2 Using Districts as the Unit of Choice in Spatial Models

Microfounded spatial models must specify the set of locations that agents can choose between. There is little guidance on what the correct resolution is—it depends on how agents perceive the city and the relationships that determine endogenous model quantities like local amenities or productivity. Moreover, the resolution also determines the dimensionality of the model, which can create a variety of technical difficulties. In many cases, districts may be both more substantively correct and more practically convenient than tracts.

First, consider models of neighborhood choice. On the substantive side, it is unlikely that searchers in a large city know detailed characteristics of each tract (Ferreira and Wong 2020). Districts could capture either a coarser conceptualization of the city or a multi-step search process in which agents first choose a broad area and then select a specific unit within it. Moreover, tract characteristics likely do not fully capture local amenities, and the watershed that determines endogenous amenity levels is probably larger than a tract—Almagro and Domínguez-Iino (2025), for example, models amenities as endogenous to a significantly broader area. Districts could be used to measure amenities or to specify the endogenous amenity production function. On the practical side, these models are often computationally intensive. Districts reduce the dimensionality significantly and eliminate the need to aggregate tracts arbitrarily or use PUMAs (Bayer et al. 2016; Calder-Wang 2021).

Second, consider quantitative spatial models (Redding and Rossi-Hansberg 2017), in which the tract-tract commuting matrix is frequently a crucial estimation moment. Dingel and Tintelnot (2026) show that this matrix is often very sparse, creating problems related to granularity. For example, only 15 percent of the 4.6 million possible tract-tract combinations in New York City contained at least one commuter in the 2010 LODES data, and more than half of that 15 percent

[21] Chetty and Hendren (2018) provide estimates for all counties with population over 10,000, and the median and 75th percentile of county population are 26,000 and 67,000, respectively. Districts' minimum population is 12,000, and the median and 75th percentile are 39,000 and 54,000. Of course, estimates in very large counties will be quite precise, but these are exactly the cases where county-level estimates are least informative about neighborhoods.

contain only a single commuter. Using districts as the unit of analysis largely solves this problem without the need for sophisticated adjustments. There are only 29,929 possible district-district combinations in New York City, and 98.3 percent contain at least one commuter in the 2010 LODES data (with over 91.5 percent of those containing three or more commuters).

7 Conclusion

We develop a new neighborhood geography using revealed preference. To make our "districts" as easy as possible to use, we construct them as discrete combinations of census tracts that are nested within counties, filling a large gap in the existing set of U.S. substate geographies. Our methodology can easily be implemented in other settings with granular migration data, including many countries' national statistical systems and linked historical complete count censuses in the U.S.. While we do not claim that districts are strictly better than census tracts or other neighborhood geographies, they have a number of attractive features and offer advantages in many applications.

Our results also point to several areas for further research. First, the extent of racial and economic clustering within districts strongly suggests that many local attributes (or perceptions of attributes) are determined at a larger scale than the tract. Examining the relationship between specific neighborhood features and local demographic composition at different scales could help open up the black boxes of neighborhood effects and endogenous amenity production. Second, the importance of physical barriers in defining neighborhoods is striking. Do these barriers represent equilibrium focal points or true breakpoints in neighborhood attributes? How do they influence neighborhood sorting over long time periods, and how could they be incorporated into spatial models of neighborhood change? Finally, people may perceive neighborhoods at relatively broad scales or make location decisions based on coarse mental maps. How might this affect common models of neighborhood choice or neighborhood equilibrium concepts? These questions represent only the starting point of what we might learn from increased curiosity about neighborhood definition and scale.

References

Aliprantis, D., Carroll, D. R., and Young, E. R. (2022). What explains neighborhood sorting by income and race? *Journal of Urban Economics*, page 103508.

Almagro, M. and Domínguez-Iino, T. (2025). Location sorting and endogenous amenities: Evidence from amsterdam. *Econometrica*, 93(3):1031–1071.

Avery, C. N., Glickman, M. E., Hoxby, C. M., and Metrick, A. (2013). A revealed preference ranking of us colleges and universities. *Quarterly Journal of Economics*, 128(1):425–467.

Bagagli, S. (2023). The (express) way to segregation: Evidence from chicago. *Job Market Paper*.

Bartik, A., Gupta, A., and Milo, D. (2025). The costs of housing regulation: Evidence from generative regulatory measurement. *Available at SSRN 4627587*.

Baum-Snow, N. and Han, L. (2024). The microgeography of housing supply. *Journal of Political Economy*, 132(6):1897–1946.

Bayer, P., Casey, M., McCartney, W. B., Orellana-Li, J., and Zhang, C. (2025). Distinguishing causes of neighborhood racial change: A nearest-neighbor design. *American Economic Review*, 115(11):3999–4039.

Bayer, P., Charles, K. K., and Park, J. (2021). Separate and unequal: Race and the geography of the American housing market. *Unpublished manuscript*.

Bayer, P., McMillan, R., Murphy, A., and Timmins, C. (2016). A dynamic model of demand for houses and neighborhoods. *Econometrica*, 84(3):893–942.

Bayer, P., Ross, S. L., and Topa, G. (2008). Place of work and place of residence: Informal hiring networks and labor market outcomes. *Journal of Political Economy*, 116(6):1150–1196.

Bergman, P., Chetty, R., DeLuca, S., Hendren, N., Katz, L. F., and Palmer, C. (2024). Creating moves to opportunity: Experimental evidence on barriers to neighborhood choice. *American Economic Review*, 114(5):1281–1337.

Billings, S. B., Hoekstra, M., and Rotger, G. P. (2022). The scale and nature of neighborhood effects on children: Evidence from a danish social housing experiment. Technical report, National Bureau of Economic Research.

Bogin, A., Doerner, W., and Larson, W. (2019). Local house price dynamics: New indices and stylized facts. *Real Estate Economics*, 47(2):365–398.

Brinkman, J. and Lin, J. (2024). Freeway revolts! the quality of life effects of highways. *Review of Economics and Statistics*, 106(5):1268–1284.

Buchanan, L. (2023). An extremely detailed guide to an extremely detailed map of New York City neighborhoods. *The New York Times*. October 29. www.nytimes.com/interactive/2023/10/29/upshot/new-york-neighborhood-guide.html.

Calder-Wang, S. (2021). The distributional impact of the sharing economy on the housing market. *Available at SSRN 3908062*.

Card, D., Mas, A., and Rothstein, J. (2008). Tipping and the dynamics of segregation. *Quarterly Journal of Economics*, 123(1):177–218.

Chetty, R., Friedman, J. N., Hendren, N., Jones, M. R., and Porter, S. R. (2025). The Opportunity Atlas: Mapping the childhood roots of social mobility. *American Economic Review*, Forthcoming.

Chetty, R. and Hendren, N. (2018). The impacts of neighborhoods on intergenerational mobility II: County-level estimates. *Quarterly Journal of Economics*, 133(1):1163–1228.

Chyn, E. and Katz, L. F. (2021). Neighborhoods matter: Assessing the evidence for place effects. *Journal of Economic Perspectives*, 35(4):197–222.

Clapp, J. M. and Wang, Y. (2006). Defining neighborhood boundaries: Are census tracts obsolete? *Journal of Urban Economics*, 59(2):259–284.

Contat, J. and Larson, W. D. (2024). A flexible method of housing price index construction using repeat-sales aggregates. *Real Estate Economics*.

Cook, C. (2025). Heterogeneous preferences for neighborhood amenities: Evidence from gps data. *Review of Economics and Statistics*, pages 1–43.

Coulton, C. J., Jennings, M. Z., and Chan, T. (2013). How big is my neighborhood? Individual and contextual effects on perceptions of neighborhood scale. *American Journal of Community Psychology*, 51:140–150.

Daepp, M. I. (2022). Small-area moving ratios and the spatial connectivity of neighborhoods: Insights from consumer credit data. *Environment and Planning B: Urban Analytics and City Science*, 49(3):1129–1146.

Davis, D., Easton, M., and Thies, S. (2025). Segregation, spillovers, and the locus of racial change. Working Paper.

Dingel, J. I. and Tintelnot, F. (2026). Spatial economics for granular settings. Technical report.

Echenique, F. and Fryer Jr, R. G. (2007). A measure of segregation based on social interactions. *Quarterly Journal of Economics*, 122(2):441–485.

Ferreira, F. V. and Wong, M. (2020). Estimating preferences for neighborhood amenities under

imperfect information. Technical report, National Bureau of Economic Research.

Fogel, J. and Modenesi, B. (2023). What is a labor market? classifying workers and jobs using network theory. *arXiv preprint arXiv:2311.00777*.

Galster, G. (2001). On the nature of neighbourhood. *Urban Studies*, 38(12):2111–2124.

Garboden, P. M. (2021). You can't get there from here: Mobility networks and the housing choice voucher program. *Journal of Planning Education and Research*, 44(3):1257–1274.

Grannis, R. (1998). The importance of trivial streets: Residential streets and residential segregation. *American Journal of Sociology*, 103(6):1530–1564.

Grannis, R. (2005). T-communities: Pedestrian street networks and residential segregation in Chicago, Los Angeles, and New York. *City & Community*, 4(3):295–321.

Guerrieri, V., Hartley, D., and Hurst, E. (2013). Endogenous gentrification and housing price dynamics. *Journal of Public Economics*, 100:45–60.

Hamilton, R. and Rae, A. (2020). Regions from the ground up: A network partitioning approach to regional delineation. *Environment and Planning B: Urban Analytics and City Science*, 47(5):775–789.

Hanson, A. and Rohlin, S. (2013). Do spatially targeted redevelopment programs spillover? *Regional Science and Urban Economics*, 43(1):86–100.

Hartley, D. A. and Rose, J. (2023). Blockbusting and the challenges faced by Black families in building wealth through housing in the postwar United States. FRB Chicago Working Paper.

Jacobs, J. (1961). *The Death and Life of Great American Cities*. Random House, New York.

Kawamoto, T. and Rosvall, M. (2015). Estimating the resolution limit of the map equation in community detection. *Physical Review E*, 91(1):012809.

L.A. Times (2025). Mapping L.A. https://maps.latimes.com/about/index.html.

Lee, B. A., Reardon, S. F., Firebaugh, G., Farrell, C. R., Matthews, S. A., and O'Sullivan, D. (2008). Beyond the census tract: Patterns and determinants of racial segregation at multiple geographic scales. *American Sociological Review*, 73(5):766–791.

Logan, J. R. (2011). Separate and unequal: The neighborhood gap for blacks, Hispanics and Asians in metropolitan America. *Project US2010 report*, pages 1–22.

Logan, J. R., Xu, Z., and Stults, B. J. (2014). Interpolating US decennial census tract data from as early as 1970 to 2010: A longitudinal tract database. *Professional Geographer*, 66(3):412–420.

Logan, T. D. and Parman, J. M. (2017). The national rise in residential segregation. *Journal of*

Economic History, 77(1):127–170.

Luttmer, E. F. (2005). Neighbors as negatives: Relative earnings and well-being. *Quarterly Journal of Economics*, 120(3):963–1002.

Mogstad, M., Romano, J. P., Shaikh, A. M., and Wilhelm, D. (2024). Inference for ranks with applications to mobility across neighbourhoods and academic achievement across countries. *Review of Economic Studies*, 91(1):476–518.

Monarrez, T. and Schönholzer, D. (2023). Dividing lines: racial segregation across local government boundaries. *Journal of Economic Literature*, 61(3):863–887.

Nimczik, J. S. (2023). Job mobility networks and endogenous labor markets.

Pebley, A. and Sastry, N. (2009). Our place: Perceived neighborhood size and names in Los Angeles. *UCLA CCPR Population Working Papers*.

Perry, C. A. (1929). *Neighborhood and Community Planning*, volume 7 of *Regional Survey of New York and Its Environs*. Regional Plan of New York and its Environs, New York.

Reardon, S. F., Fox, L., and Townsend, J. (2015). Neighborhood income composition by household race and income, 1990–2009. *Annals of the American Academy of Political and Social Science*, 660(1):78–97.

Reardon, S. F., Matthews, S. A., O'sullivan, D., Lee, B. A., Firebaugh, G., Farrell, C. R., and Bischoff, K. (2008). The geographic scale of metropolitan racial segregation. *Demography*, 45(3):489–514.

Redding, S. J. and Rossi-Hansberg, E. (2017). Quantitative spatial economics. *Annual Review of Economics*, 9:21–58.

Rosvall, M., Axelsson, D., and Bergstrom, C. T. (2009). The map equation. *European Physical Journal Special Topics*, 178(1):13–23.

Rosvall, M. and Bergstrom, C. T. (2008). Maps of random walks on complex networks reveal community structure. *Proceedings of the National Academy of Sciences*, 105(4):1118–1123.

Sampson, R. J. and Sharkey, P. (2008). Neighborhood selection and the social reproduction of concentrated racial inequality. *Demography*, 45(1):1–29.

Schelling, T. C. (1971). Dynamic models of segregation. *Journal of Mathematical Sociology*, 1(2):143–186.

Schmutte, I. M. (2014). Free to move? A network analytic approach for learning the limits to job mobility. *Labour Economics*, 29:49–61.

Schönholzer, D. (2023). Measuring preferences for local governments.

Siordia, C. and Wunneburger, D. F. (2013). Contiguity principle for geographic units: Evidence on the quantity, degree, and location of Public Use Microdata Area (PUMA) fragmentation. *Human Geographies*, 7(2):5.

Sorkin, I. (2018). Ranking firms using revealed preference. *Quarterly Journal of Economics*, 133(3):1331–1393.

Spielman, S. E. and Logan, J. R. (2013). Using high-resolution population data to identify neighborhoods and establish their boundaries. *Annals of the Association of American Geographers*, 103(1):67–84.

Suttles, G. (1972). *The Social Construction of Communities*. University of Chicago Press.

Tach, L., Pendall, R., and Derian, A. (2016). Income mixing across scales: Rationale, trends, policies, practice, and research for more inclusive neighborhoods and metropolitan areas.

Talen, E., Wileden, L., and Bae, C. (2025). Normative neighborhoods: Does theory match perception? Manuscript, University of Chicago and University of Connecticut.

Voorheis, J. L., Colmer, J., Houghton, K., Lyubich, E., Munro, M., Scalera, C., and Withrow, J. (2023). Building the prototype census environmental impacts frame. Technical report, National Bureau of Economic Research.

Wang, S. (2025). Are highways conduits or barriers for urban travelers? a welfare analysis using smartphone data.

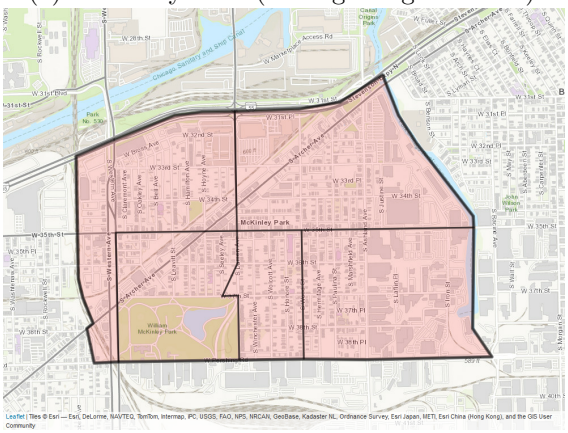
White, M. J. (1983). The measurement of spatial segregation. *American Journal of Sociology*, 88(5):1008–1018.

Woodruff, A. (2017). Official unofficial neighborhoods, 2017. *Bostonography*. April 4. https://bostonography.com/2017/official-unofficial-neighborhoods-2017.

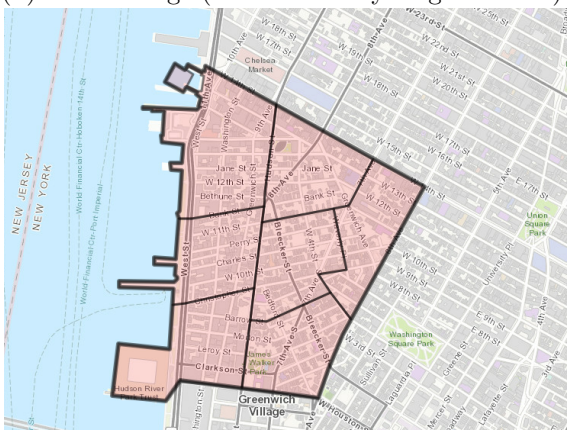
Ziff, A. (2025). Beyond the local impacts of place-based policies: Spillovers through latent housing markets. Working Paper.

Figure 1: Example Districts

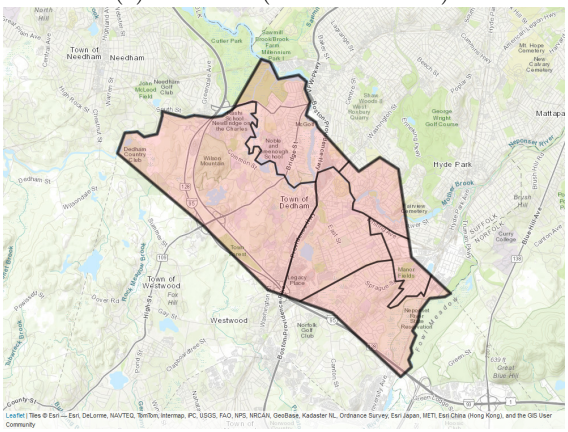
(a) McKinley Park (Chicago neighborhood)



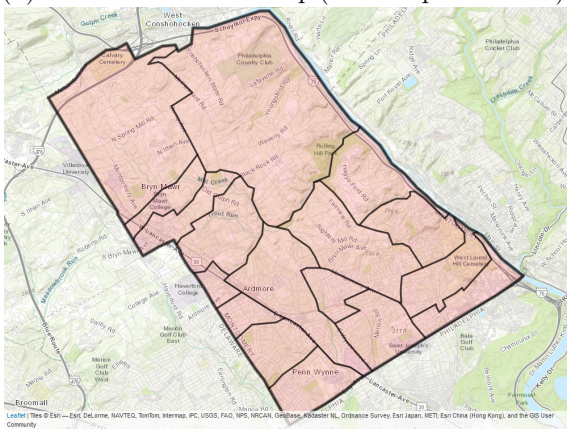
(b) West Village (New York City neighborhood)



(c) Dedham (Boston suburb)



(d) Lower Merion Township (Philadelphia suburb)



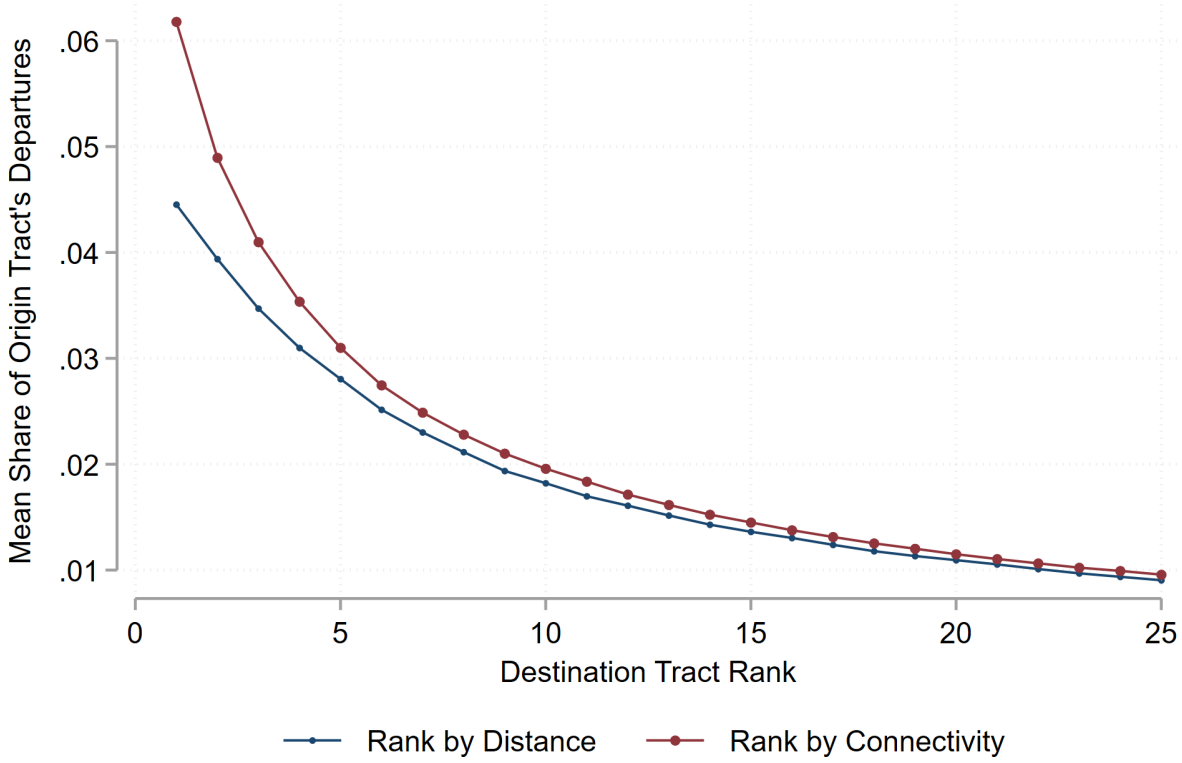
Note: This figure provides illustrative examples of the district neighborhood concept. The district of interest is highlighted in pink, and the tracts that compose it are outlined in black.

The upper two examples are cases where district boundaries reflect popular perceptions of neighborhoods within cities. Panel A shows a district that aligns closely with the McKinley Park neighborhood on the southwest side of Chicago. Its population is about 16,000, and its borders are formed by a combination of industrial areas, waterways, and railroad tracks. The district in Panel B matches popular perceptions of the West Village neighborhood in lower Manhattan (population 35,000), whose borders are formed by large streets and the Hudson River.

In the lower examples, districts align with important municipal boundaries. Panel C shows a district that coincides nearly exactly with the town of Dedham, Massachussetts (population 25,000), and Panel D shows a district that is equivalent to Lower Merion Township, Pennsylvania (population 65,000).

Sources: District boundaries and U.S. Census TIGER/Line Shapefiles. Basemap by Esri via Leaflet.

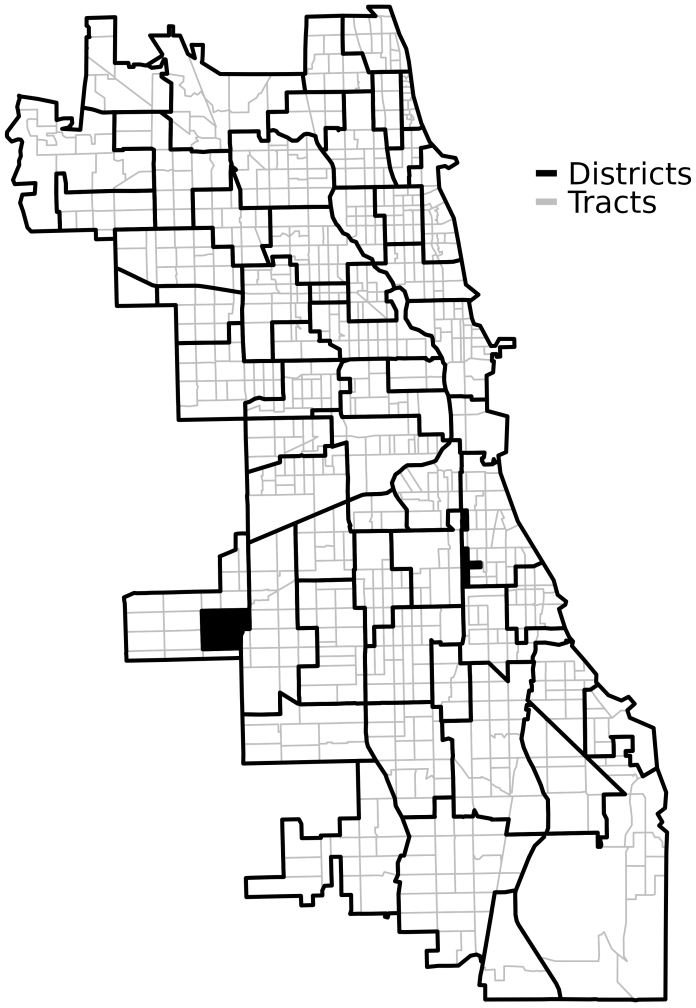
Figure 2: Share of Within-County Departures to Closest Neighbors and Strongest Connections



Note: This figure shows the share of the average census tract's out-migrants who move to each of the tract's 50 nearest neighbors and the tract's 50 strongest migration connections. The migration connection strength between two tracts is defined in the same way as in the district estimation algorithm. We restrict to tract pairs that are in the same county, and we include only tracts in counties that have at least 75 tracts and are in a CBSA with population over 250,000.

Sources: Authors' calculations using EIF data.

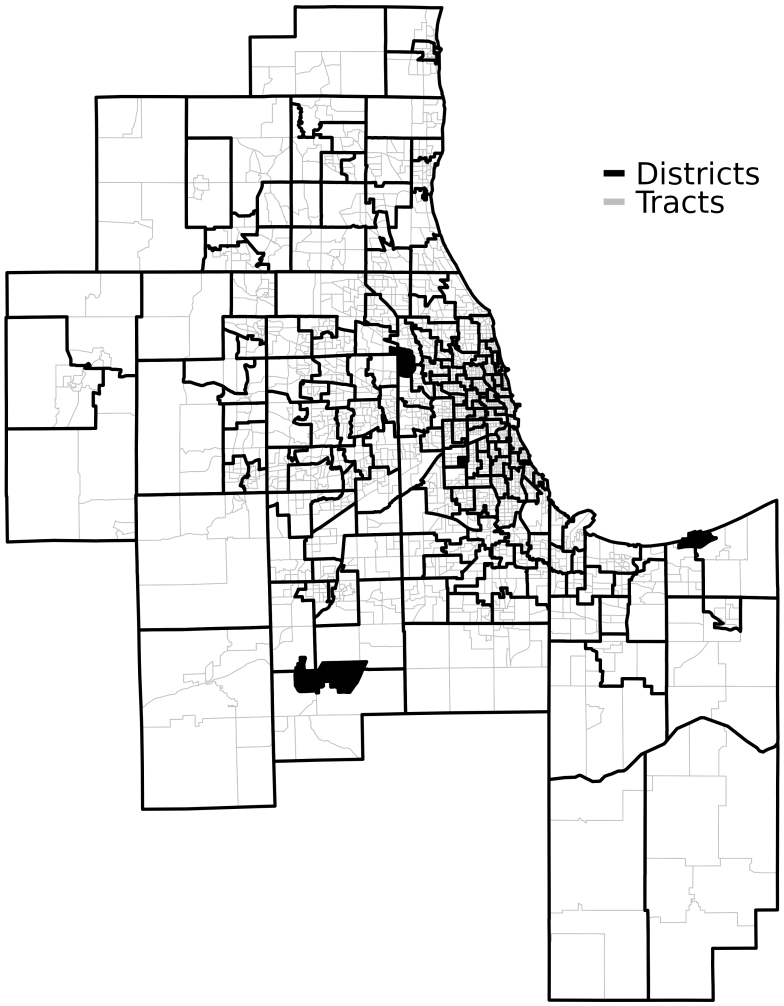
Figure 3: District Boundaries in the City of Chicago



Note: This figure shows district boundaries for the city of Chicago. Thick lines represent district boundaries, while thin lines show census tracts. Tracts with black fill have population below 500 and are not assigned to a district.

Sources: District boundaries and U.S. Census TIGER/Line Shapefiles.

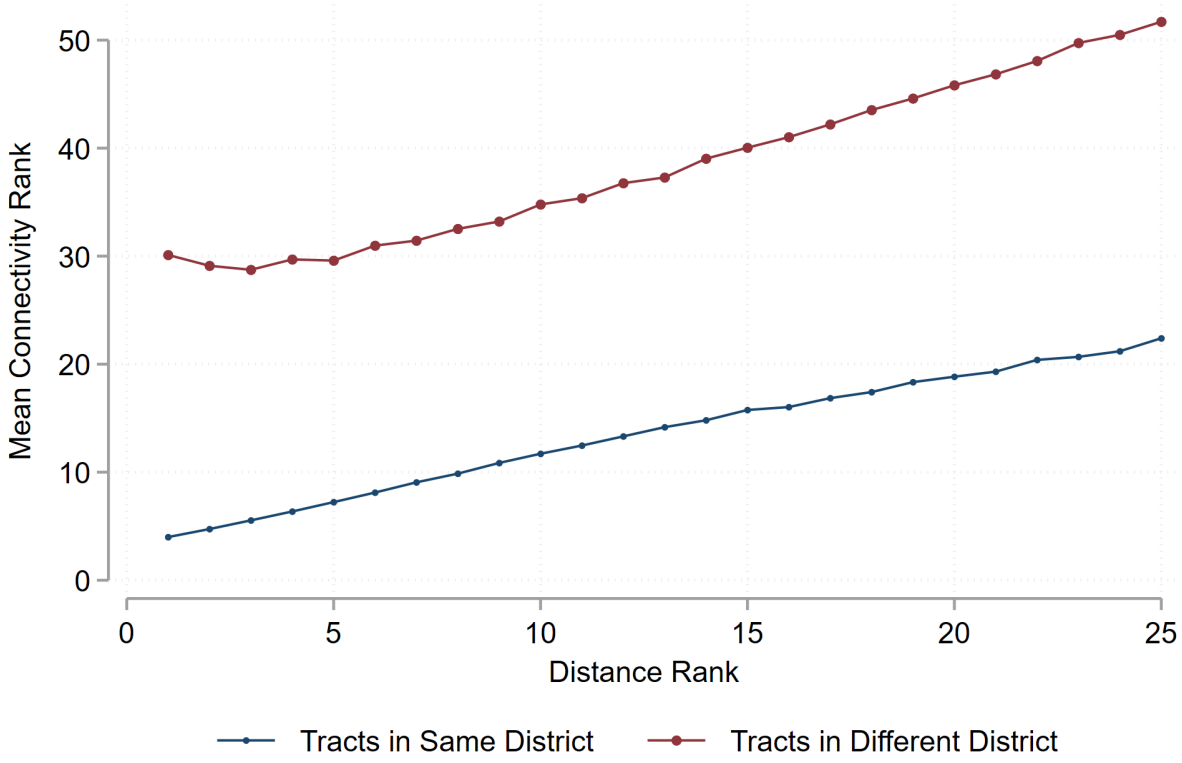
Figure 4: District Boundaries in Chicago CBSA



Note: This figure shows district boundaries in the Chicago CBSA. Thick lines represent district boundaries, while thin lines show census tracts. Tracts with black fill have population below 500 and are not assigned to a district.

Sources: District boundaries and U.S. Census TIGER/Line Shapefiles.

Figure 5: Tract Migration Connectivity Versus Distance, Within and Across Districts

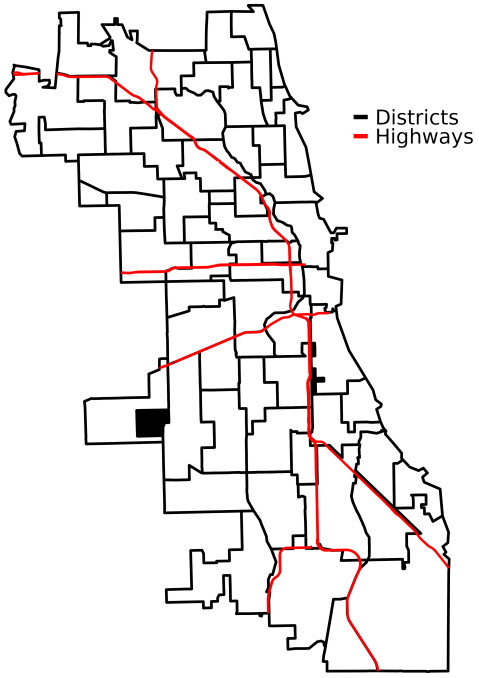


Note: This figure shows the average migration connectivity rank between a tract and its 25 nearest neighbors, stratified by whether the neighboring tract is in the same district as the index tract. For example, the leftmost dot of the lower line shows that when a tract's closest neighbor is in the same district, it is the index tract's fourth-strongest migration connection on average. The leftmost dot of the upper line shows that when the closest neighbor is instead a different district, it is the 30th strongest connection. The migration connection strength between two tracts is defined in the same way as in the district estimation algorithm. We restrict to tract pairs that are in the same county, and we include only tracts in counties that have at least 75 tracts and are in a CBSA with population over 250,000.

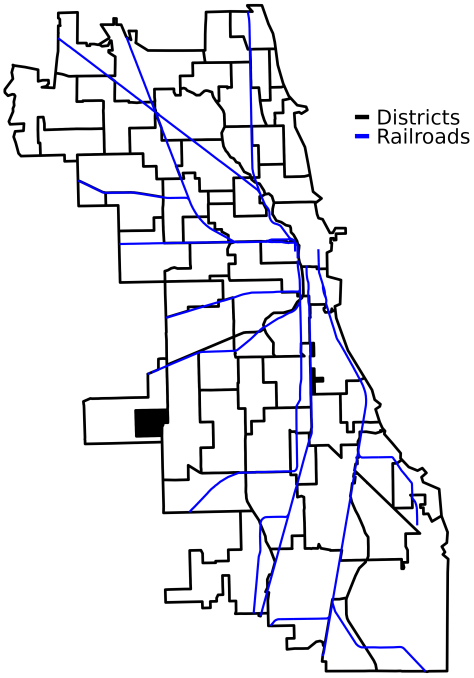
Sources: Authors' calculations using district boundaries and EIF data.

Figure 6: Physical Barriers and District Boundaries in Chicago

Panel A: Interstate Highways



Panel B: Railroad Tracks

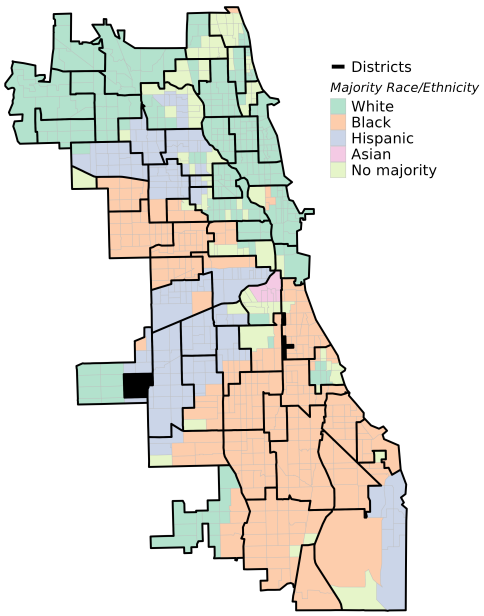


Note: This figure plots Chicago districts against interstate highways (left panel) and major railroad lines (right panel, Chicago Transit Authority lines excluded). Tracts with black fill have population below 500 and are not assigned to a district.

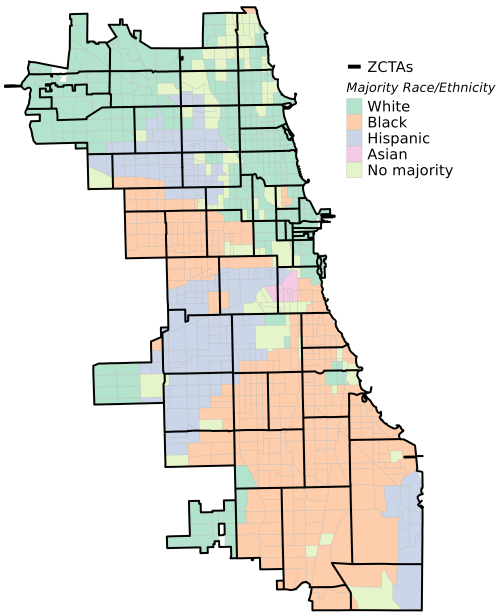
Sources: District boundaries, City of Chicago Data Portal Railroads and Major Streets shapefiles.

Figure 7: Tract Racial Majority versus District and Zip Code Boundaries

Panel A: Districts



Panel B: ZCTAs



Note: This figure overlays the boundaries of districts (left panel) and Zip Code Tabulation Areas (right panel) on a map of tract racial majorities in Chicago. Tracts with black fill have been dropped from the district algorithm because their population is below 500.

Sources: District boundaries and 2010 decennial census populations as compiled in the LTDB.

Table 1: Substate Geographies in the United States

	County	**PUMA**	**ZCTA**	**Tract**	**Block group**	**Block**	**District**
Typical population range	73–9,974,000	100,000–200,000	0–115,000	1,200–8,000	600–3,000	30–600	20,000–80,000
Nesting	States > counties > tracts	States > PUMAs > tracts	States > ZCTAs > blocks	Counties > tracts > block group	Tracts > block group > blocks	Block group > blocks	Counties > districts > tracts
Contiguous	No	No	No	Yes	Yes	Yes	Yes
Initial purpose	Local government unit	Microdata location identifier	Mail delivery	Statistical tabulation	Statistical tabulation	Statistical tabulation	Neighbor-hood unit
Available Census data	Microdata w/county identifiers	Microdata w/PUMA identifiers	Aggregates	Aggregates	Limited aggregates	Minimal aggregates	Aggregates
Origin	State governments	Census Bureau	U.S.P.S. /Census	Census Bureau	Census Bureau	Census Bureau	This paper

Note: This table describes features of several popular substate geographies. PUMAs are Public Use Microdata Areas, and ZCTAs are Zip Code Tabulation Areas. ZCTAs are collections of census blocks that match postal zip codes as closely as possible. Since counties and ZCTAs do not have a target size range, we instead report the extremes of the population distribution.

Sources: The Census Bureau's Geographic Areas Reference Manual, authors' calculations using 2010-2014 ACS.

Table 2: District Size Distribution

	P10	P25	Median	P75	P90	Mean	*N*
Full sample							
Population	20,000	27,400	38,500	53,500	69,800	42,300	4,762
Tract count	5	6	9	12	16	9.7	4,762
Land area (sq. miles)	3.1	6.9	16.8	56.1	197.2	86.7	4,762
Top quartile town pop. density							
Population	21,100	29,100	40,800	56,100	73,400	44,400	1,149
Tract count	6	7	10	14	18	11.2	1,149
Land area (sq. miles)	1.0	2.2	4.3	7.7	13.4	7.2	1,149
Lower quartiles town pop. density							
Population	19,700	27,400	38,500	53,500	69,800	42,300	3,613
Tract count	5	6	9	12	16	9.7	3,613
Land area (sq. miles)	3.1	6.9	16.8	56.1	197.2	86.7	3,613

Note: This table shows the distribution of three measures of district size. In the bottom two panels, districts are separated according to the population density quartile of their census place. If a district spans multiple census places, it classified according to the densest.

Sources: Authors' calculations using district boundaries, tract land areas from U.S. Census TIGER/Line Shapefiles, and 2010 decennial census populations and Census Place designations as compiled in the LTDB.

Table 3: Alignment Between District Borders, Physical Barriers, and Administrative Boundaries

Type of adjacent tract pair	Share that contain district border	Obs.
All tract pairs		
Not separated by highway	0.301	105,800
Separated by highway	0.483	8,000
Low pop. density		
Not separated by highway	0.247	53,200
Separated by highway	0.326	3,900
High pop. density		
Not separated by highway	0.356	52,500
Separated by highway	0.634	4,100
Tracts in same town	0.253	79,600
Tracts in different towns	0.562	24,800
Tracts in same school district	0.275	102,900
Tracts in different school districts	0.676	10,900
Same school district		
Tracts in same town	0.223	65,900
Tracts in different towns	0.481	15,700
Different school districts		
Tracts in same town	0.410	1,300
Tracts in different towns	0.737	4,800

Note: This table shows the alignment between district borders, municipal borders, school district borders, and interstate highways. An observation is a pair of adjacent census tracts, and, because districts are estimated within counties, only tract pairs within the same county are included. We define highways as limited-access interstate highways or state-managed roads with ramps and interchanges. We say that a tract border aligns with a highway if at least 50 percent of the border (drawn at a 20-foot width), is covered by highway (drawn at a 60-foot width per Department of Transportation design standards for interstate highways). We define a tract pair as high-density if the densest district with which it intersects is above the median district density. We define a pair of tracts as in the same town if they are both classified into the same Census Place, and we include only pairs in which both tracts are part of a Census Place when considering this variable. Finally, we define the alignment of school district and tract boundaries in a similar way as highways, and we include boundaries between all levels of school districts (i.e., elementary, secondary, and unified). Note that the methodology for constructing the highway and school district indicators cannot be applied to tract pairs that intersect only at a single point, so “queen" tract pairs are excluded from the sample when considering these variables.

Sources: Authors’ calculations using district boundaries, 2010 decennial census populations and Census Place designations as compiled in the LTDB, and from U.S. Census: TIGER/Line Shapefiles, 2019 Primary Roads Shapefile, and 2019 School District Boundary Shapefiles.

Table 4: Share of Districts, Zip Codes, and Buffer Regions Containing a Demographic Discontinuity

	White share discontinuity			*Income discontinuity*		
Tract count	Districts	ZCTAs	Buffers	Districts	ZCTAs	Buffers
4	3.4%	5.9%	10.8%	4.0%	9.2%	11.9%
6	8.5%	11.9%	16.8%	12.4%	17.1%	17.6%
8	11.7%	20.9%	21.1%	18.3%	26.0%	22.3%
10	13.8%	26.2%	27.0%	23.6%	32.4%	28.9%
12	21.1%	27.1%	37.1%	33.5%	37.8%	38.5%

Note: This table shows the likelihood that different spatial aggregations contain a large demographic discontinuity. A large discontinuity is defined as a pair of adjacent census tracts that are in the top 5 percent of all adjacent tract pairs in the difference in their white share (left columns) or log median household income (right columns). A buffer region is a census tract plus all tracts that border it. Units are stratified according to how many census tracts they contain, with tracts being assigned to the ZCTA that contains the highest percentage of their population.

Sources: Authors' calculations using district boundaries, U.S. Census TIGER/Line Shapefiles, and demographic information compiled from the 2010 decennial census (percent white) and 2008-2012 American Community Survey aggregates (household income) in the LTDB.

Table 5: R^2 from Regressions of Tract Characteristics on District Fixed Effects

	Percent white	Percent Black	Percent Hispanic	Percent poor	Log median income
R^2	0.875	0.873	0.887	0.648	0.696
Tract obs.	46,100	46,100	46,100	46,100	46,100

Note: This table shows the amount of variation in tract characteristics that can be explained by district fixed effects. Each column represents a separate regression with the outcome variable indicated in the heading.

Sources: Authors' calculations using district boundaries and demographic information compiled from the 2010 decennial census (racial composition) and 2008-2012 American Community Survey aggregates (poverty rate, household income) in the LTDB.

Table 6: District Characteristics Versus Tract Racial Majority, Conditional on Tract Characteristics

	District poverty	Log district income	District college share
1(maj. Black tract)	0.067	-0.202	-0.039
	(0.0012)	(0.0045)	(0.0017)
1(maj. Hispanic tract)	0.045	-0.163	-0.047
	(0.0013)	(0.0046)	(0.0018)
1(other maj. minority tract)	0.023	-0.078	-0.020
	(0.00096)	(0.0036)	(0.0014)
N	46,100	46,100	46,100
R^2	0.646	0.714	0.769
Mean dependent variable	0.148	10.96	0.313
CBSA fixed effects	Y	Y	Y
Tract controls	Y	Y	Y

Note: This table shows results from regressing characteristics of the district surrounding a tract (excluding the index tract itself) on indicator variables for tract racial majority. Each column represents a separate regression with the outcome indicated in the heading. We control for tract poverty rate, college share, log median income, employment rate, and owner occupancy rate. All regressions include CBSA fixed effects. The omitted racial majority category is majority white.

Sources: Authors' calculations using district boundaries and demographic information compiled from the 2010 decennial census (racial composition and owner occupancy rate) and 2008-2012 American Community Survey aggregates (all other variables) in the LTDB.

Table 7: Aggregating Opportunity Atlas Estimates to District Level

	Original tract estimates	District aggregation
Mean S.E. (percentiles)	2.08	0.67
Signal variance explained	1	0.827
MAPE (%) in hold-out sample		
All children	7.8	7.2
White children	11	9.3
Black children	14.8	10.9
Hispanic children	14.8	10.8

Note: This table shows properties of the district aggregation of the Opportunity Atlas's tract-level estimates of intergenerational mobility for children born to parents at the 25th income percentile. We construct district aggregations of the tract estimates and their standard errors as simple population-weighted averages. The mean S.E. row reports the mean tract standard error as reported in Chetty et al. (2025) as well as the mean from our district aggregations. Signal variance explained is computed following Chetty et al. (2025). Mean Absolute Prediction Error (MAPE) in a hold-out sample is computed in the following way. First, we define the hold-out sample as the 1984-1989 birth cohorts (versus the 1978-1983 cohort used for the Opportunity Atlas's baseline estimates). We then predict tract-level mobility estimates in this hold-out sample (which are provided by the Opportunity Atlas) as either the tract's mobility estimate in the baseline sample or the average baseline mobility estimate in the tract's district. We then calculate MAPE as the mean absolute percentage deviation between the hold-out sample estimate and the predicted value.

Sources: Authors' calculations using district boundaries and Opportunity Atlas intergenerational mobility estimates.

A Online Appendix

A.1 Model Performance Metrics

This section describes the performance metrics used to choose the connection strength metric in the network representation of the tract-tract transition data. First, we precisely define the subsample stability and migratory disconnection measures.

Subsample stability. If a district is too small, we expect to see borders in places that are not actually very important determinants of migration decisions, likely because the algorithm is overfitting noise in the data. To assess overfitting, we run each specification on three ten-percent subsamples of the EIF and compute the similarity between the results from each subsample to the results from running the same specification on the full sample. We operationalize similarity as the share of adjacent tract pairs that have the same district border status in the subsample and full sample estimation.

Intra-district disconnection. If a district is too large, it should contain subclusters of tracts that have distinct migration patterns. For instance, the specification using raw migration counts generates very large districts, within which individual neighborhoods display identifiable and distinct migration flows. There is no formal definition of a subcluster, so we develop our own metric by combining k-nearest neighbor and graph reachability techniques. For every district, we create a matrix akin to a transition matrix, but that is populated with ones if tract j is among tract i's five strongest connections and zeros otherwise. We raise this matrix to the 10th power and say that a district is disconnected if there is a cell of the exponentiated matrix that is equal to zero. The intuition is that a zero cell indicates that there is a pair of tracts in the district that cannot be connected on a ten-step path built with only strong links.

Appendix Table A.4 summarizes the performance of 10 specifications on our performance metrics. The first four rows represent specifications with the connection weight from Tract A to Tract B set equal to the raw number of moves, raised to powers from one to four. For example, with exponent one, this sets the probability that a random walker moves from Tract A to Tract B equal to the observed share of departures from A who moved to B. Exponent three performs the best of these specifications, with 77 percent of districts falling in the desired size range and the highest stability across subsamples. However, all have a wide range of sizes, leading a large share of districts to fall outside the desired range even when the mean size is on target. This is likely because migration counts are heavily influenced by tract population and mobility rates, introducing noise that is not correlated with substitutability.

In the lower rows, we reduce this variance by defining connection strength based on a destination's rank among all tracts that receive migrants from the origin. Specifically, we divide the number of migrants from tract A to tract B by the total number of migrants to B, producing a measure of B's popularity among A's out-migrants that accounts for B's size. We then rank all of A's destinations according to this variable, with one being the highest. Finally, denoting this rank as R_{ab}, we define connection strength between A and B to $\frac{1}{(R_{ab}+2)^X}$ if $R_{ab} \leq C$ and 0 otherwise.

A much higher share of districts land in the desired range under parameterizations of this form than using raw migration counts. Overall, the specification with an exponent X =2 and a cutoff C =10 performs best, generating 91 percent in the target size range, the best scores on stability across subsamples, and the second best score on the disconnection metric. We use this as our primary specification.

A.2 Model Robustness

The algorithm is not particularly sensitive to small changes to the specification. Appendix Table A.5 shows the share of adjacent tract pairs that receive the same classification in the baseline versus the alternative specification indicated by the column heading. We stratify tract pairs based on the number of 10 percent subsamples in which they are placed in different districts under the baseline specification. Agreement across specifications is quite high among pairs that are nearly always or nearly never borders. Most of the disagreements between specifications occur in the middle of the subsample border count distribution. These are "fuzzy" borders that have migration connections that do not clearly indicate whether the tracts are part of the same district. This pattern suggests that small changes in parameterization will not alter the location of important borders that are true breaks in the urban landscape or randomly create borders between tracts with strong connections, but they do lead to some differences.

A.3 Data Appendix

Sample definition details: Both Core-Based Statistical Area and Census tract geographies are the 2010 vintage. To define the 100 most populous CBSAs and identify low-population tracts, we use the population totals from the 2010-2014 ACS.

Shape files and crosswalks: We take tract shape files, which are used both for mapping and for identifying the adjacent tract pairs used frequently in our analysis, from the Census Bureau. A crosswalk from tracts to zip code tabulation areas is taken from the Department of Housing and Urban Development, and a crosswalk from tracts to PUMAs is taken from the Census Bureau.

Demographic information: We take all demographic information from the Longitudinal Tract Database constructed in Logan et al. (2014). This database is constructed using decennial censuses from 1970 to 2020, supplemented by ACS information in more recent years. All variables are harmonized to 2010 vintage census tracts. We also use this database to determine the Census Place that each tract is a part of.

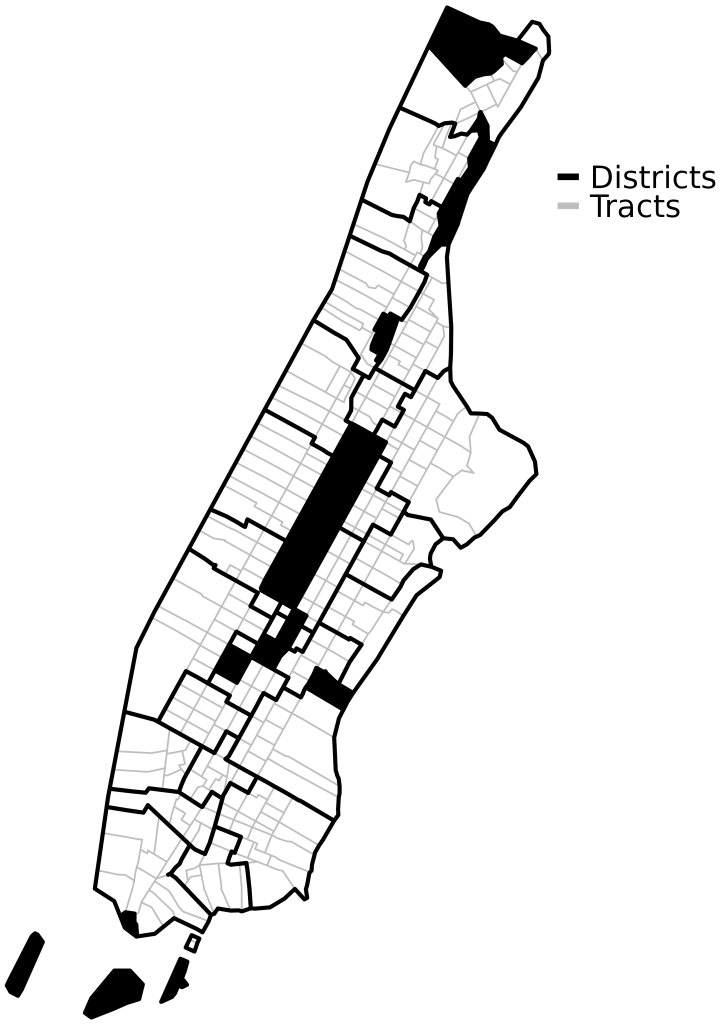
Correcting error in the FSRDC output: We incorrectly applied a crosswalk from 2020 tracts to 2010 tracts within the Federal Statistical Research Data Center and did not realize the error until after completing the irreversible disclosure process. The error leads 1.2 percent of 2010 vintage census tracts to not be assigned to a district. We corrected the error to the best of our ability with the following two steps (all taken outside the FSRDC). First, we identified 2010 vintage tracts that *did not* receive a district ID but overlapped with a 2020 vintage tract that also overlaps with a 2010 vintage tract that *did* receive a district. We assign the 2010 vintage tract that is missing a district ID to the same district as the other 2010 vintage tract that overlaps with the 2020 vintage tract. This lowers the error rate to less than 0.1 percent, essentially by accounting for places where multiple 2010 vintage tracts were merged into a single 2020 vintage tract. Second, we assign the remaining unmatched tracts to the same district as their most demographically similar bordering tract. This correction applies to 51 tracts and leaves only 14 unassigned tracts remaining. These 14 tracts were not assigned a district not because of the crosswalk error, but because they are geographic oddities that create issues with our algorithm, as described below.

Geographic oddities: In the process of creating a geography that covers most of the United States, we ran into a number of oddities that could form the basis for a rather boring bar trivia night. One recurring issue is small islands, which are often below our population threshold and have no adjacent district to be combined with. Four census tracts are literal islands (three surrounded by ocean waters, one by a Kansas City industrial agglomeration) with population well below our minimum of 12,000 and no nearby mainland. We classify these tracts as isolated and do not assign

them to a district. On the other hand, in the Seattle metro, two tracts comprise Vashon Island and have population just below 12,000—we designate this as its own district. In the New York City metro, an area of Queens is completely surrounded by water but connected to the rest of the borough by a short bridge. We classify this with the nearest district in Queens. Another issue is that some counties are discontiguous. This can create problems given that we require districts to be both contiguous and contained within a single county. If the discontiguous area is very close to the remainder of the county, we simply lump it in with the nearest district. If the area is instead far (e.g., Cohasset in Norfolk County, MA), we leave it as its own district even if slightly below the population threshold.

Appendix Figures and Tables

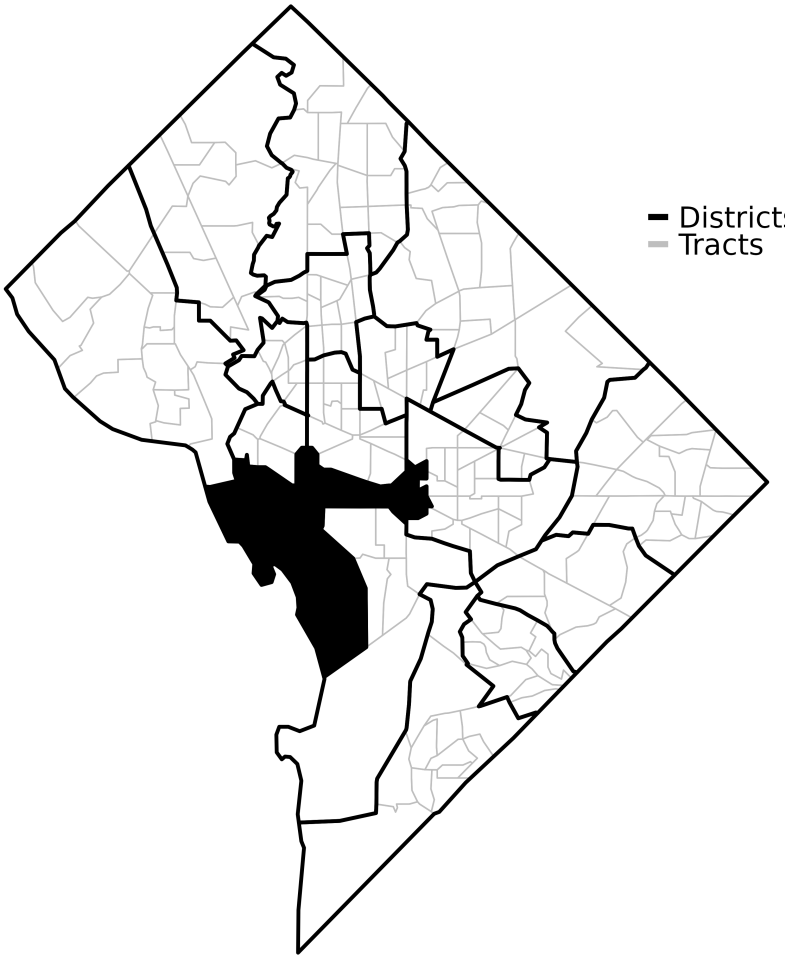
Figure A.1: District Boundaries in Manhattan



Note: This figure shows district boundaries in Manhattan. Thick lines represent district boundaries, while thin lines show census tracts. Tracts with black fill have been dropped from the algorithm because their population is below 500.

Sources: District boundaries and U.S. Census TIGER/Line Shapefiles.

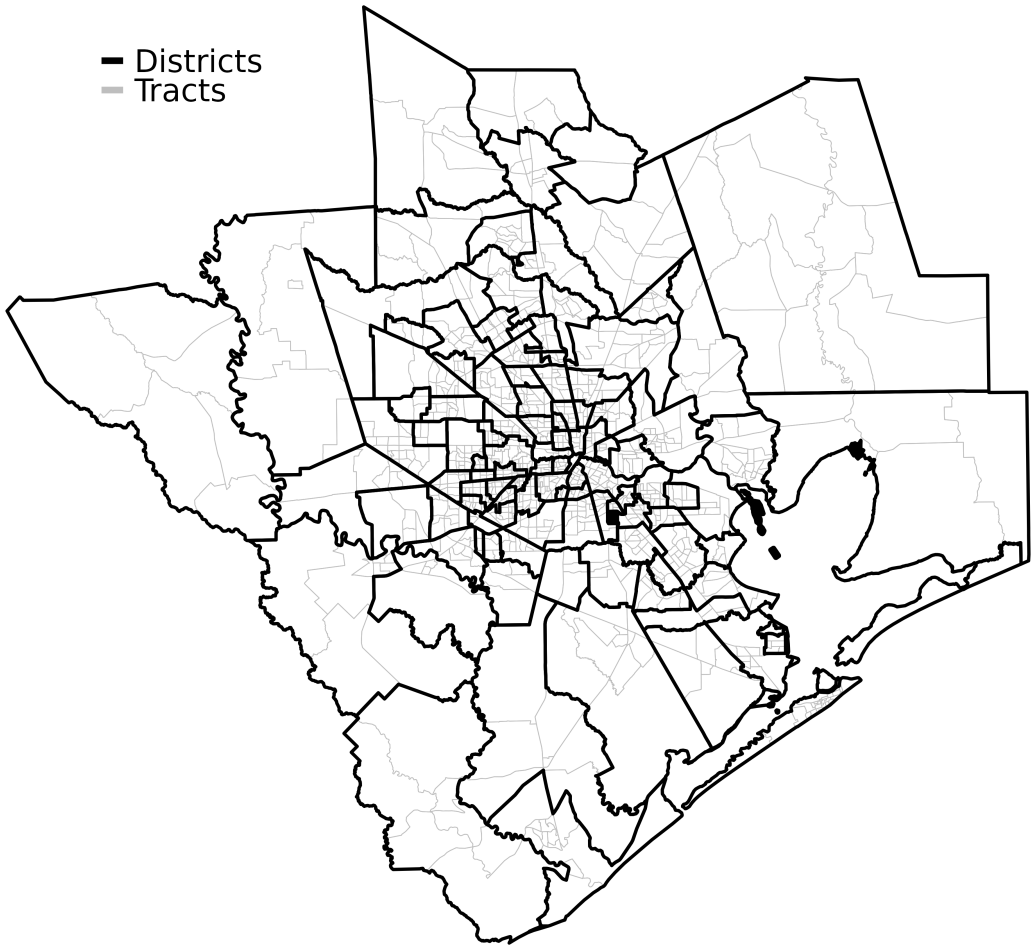
Figure A.2: District Boundaries in Washington, D.C.



Note: This figure shows district boundaries in Washington, D.C. Thick lines represent district boundaries, while thin lines show census tracts. Tracts with black fill have population below 500 and are not assigned to a district.

Sources: District boundaries and U.S. Census TIGER/Line Shapefiles.

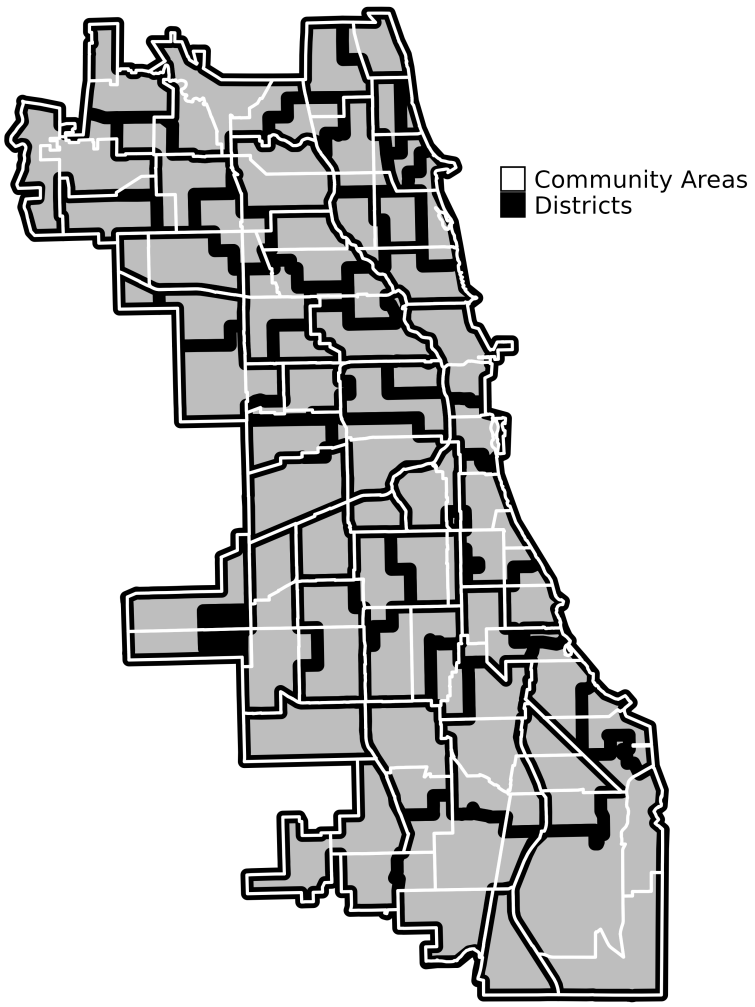
Figure A.3: District Boundaries in Houston CBSA



Note: This figure shows district boundaries in the Houston metropolitan area. Thick lines represent district boundaries, while thin lines show census tracts. Tracts with black fill have population below 500 and are not assigned to a district.

Sources: District boundaries and U.S. Census TIGER/Line Shapefiles.

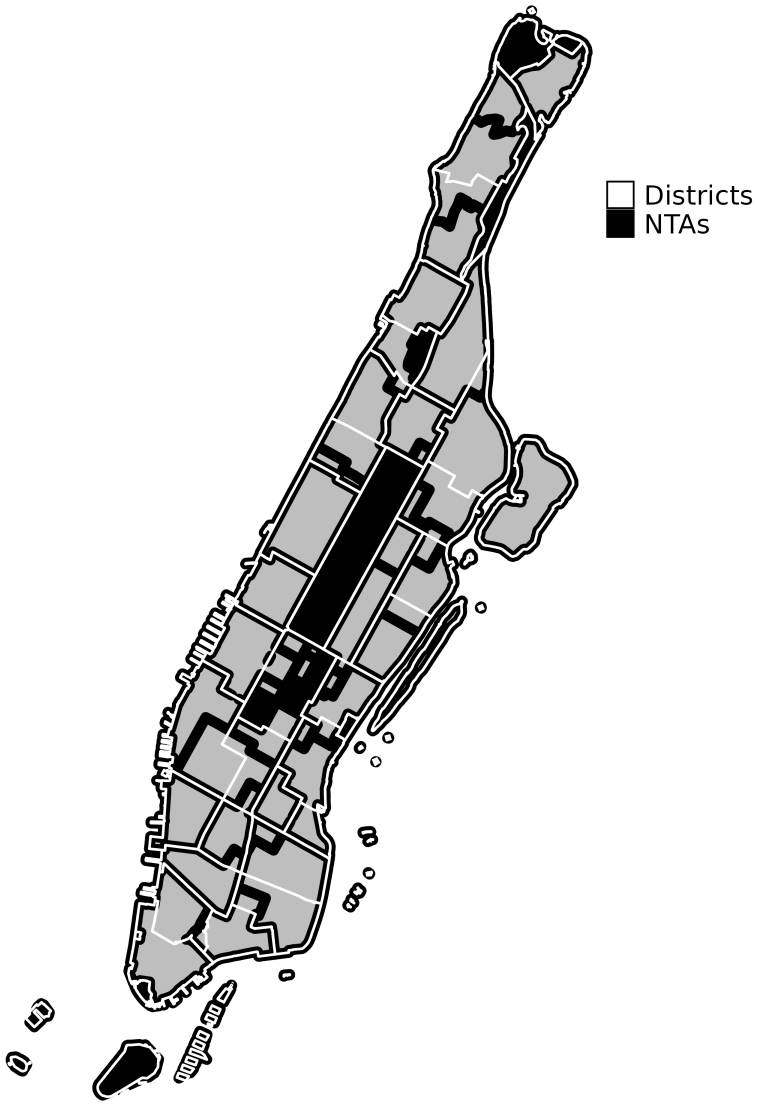
Figure A.4: Chicago Community Areas Versus Districts



Note: This figure plots Chicago districts against the city's Community Areas, a popular unofficial neighborhood definition that was created by demographers in the 1920s. Community Area boundaries are shown in white in the foreground, while the district borders are shown in the thicker black lines. Tracts with black fill have population below 500 and are not assigned to a district.

Sources: District boundaries, U.S. Census TIGER/Line Shapefiles, and City of Chicago Data Portal Community Areas Boundaries shapefile.

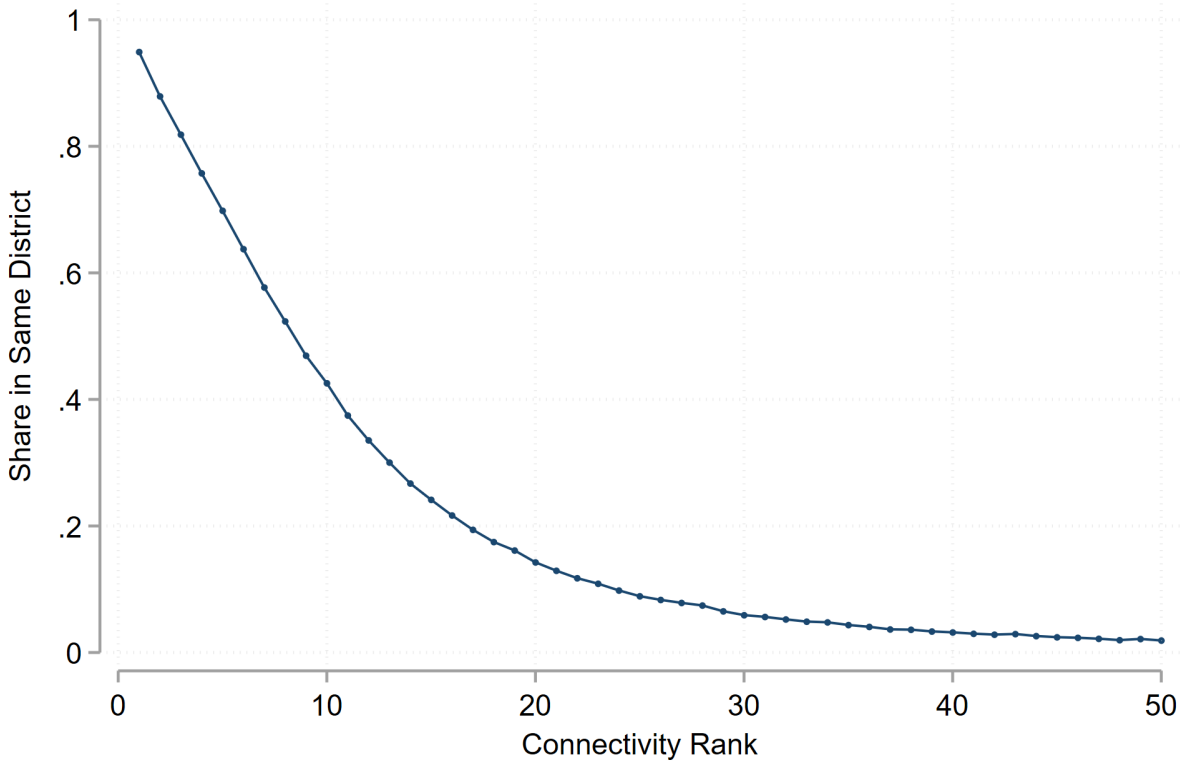
Figure A.5: Manhattan Neighborhood Tabulation Areas vs. Districts



Note: This figure plots New York City's Neighborhood Tabulation Areas against districts in Manhattan. NTA boundaries are shown in white in the foreground, while the district borders are shown in the thicker black lines. Tracts with black fill have population below 500 and are not assigned to a district.

Sources: District boundaries, U.S. Census TIGER/Line Shapefiles, and NYC Open Data 2020 Neighborhood Tabulation shapefile.

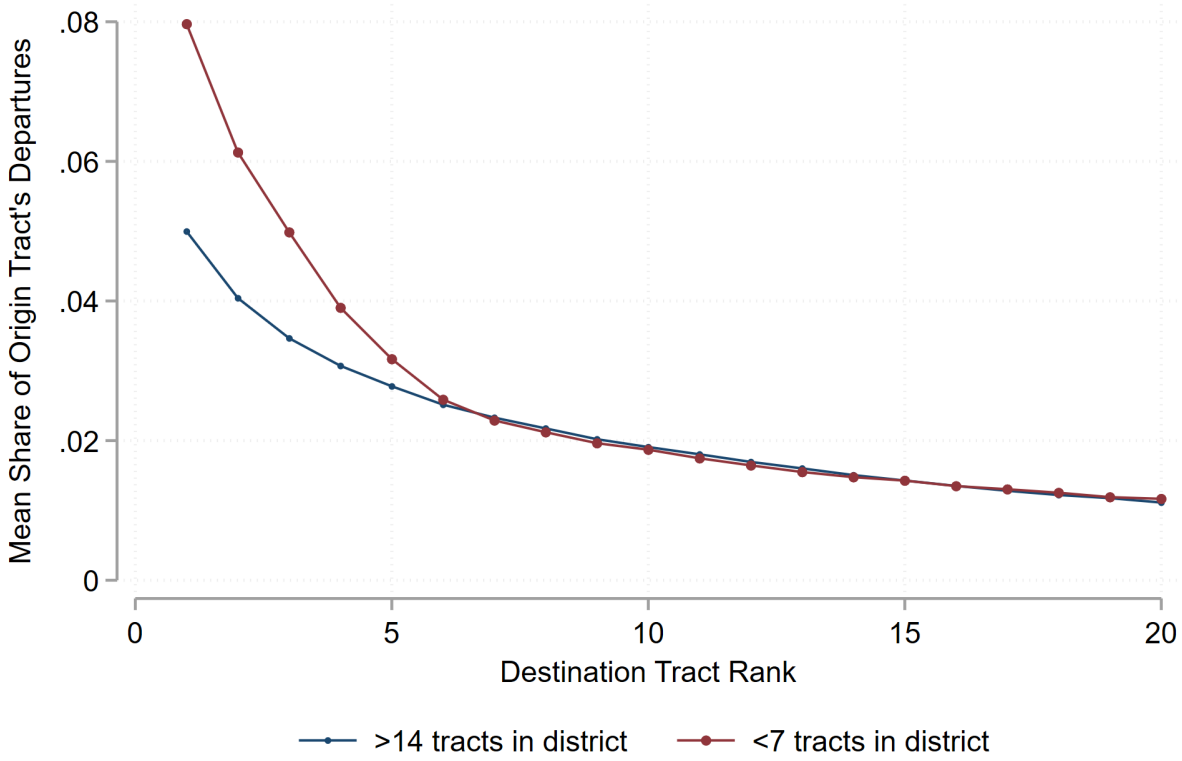
Figure A.6: Share of Strongest Connections Classified into Same District



Note: This figure shows, for each of a tract's 50 strongest migratory connections, the share that are placed in the same district as the index tract. The connection strength between two tracts is defined in the same way as in the district estimation algorithm. We restrict to tract pairs that are in the same county, and we include only tracts in counties that have at least 75 tracts and are in a CBSA with population over 250,000.

Sources: Authors' calculations using district boundaries and EIF data.

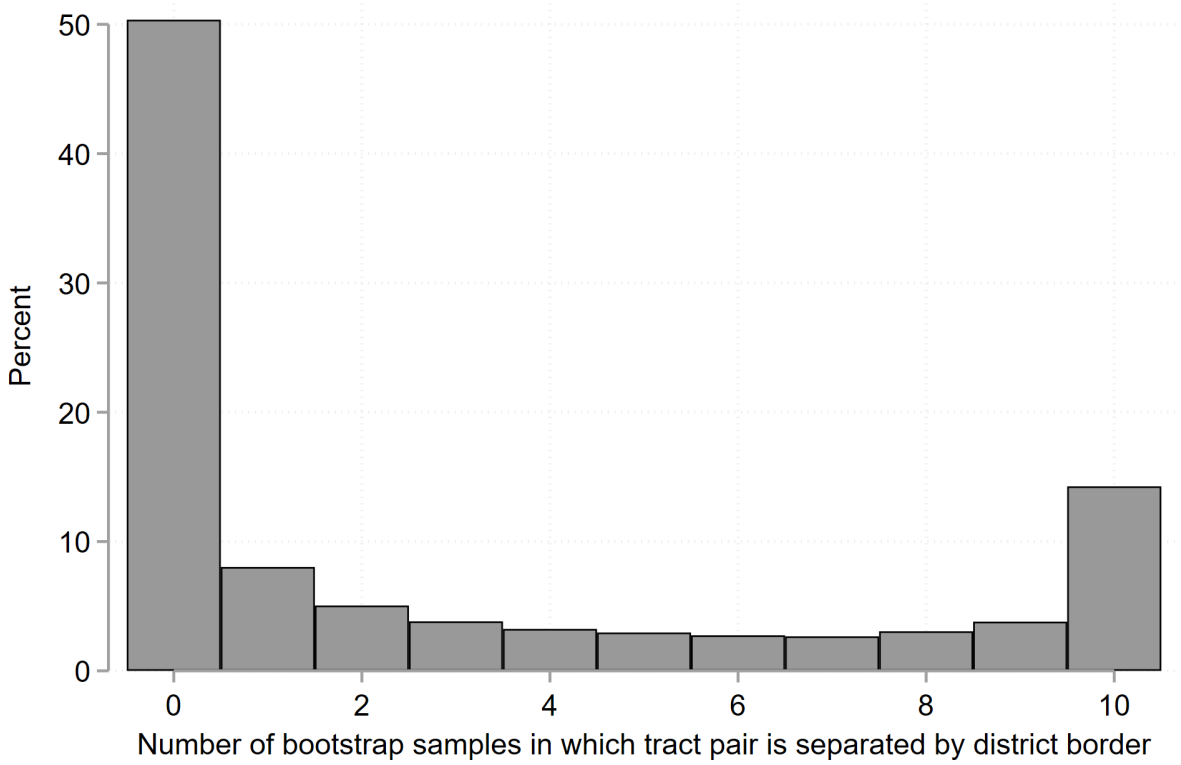
Figure A.7: Share of Departures to Strongest Connections in Large Versus Small Districts



Note: This figure shows the average share of a census tract's out-migrants who move to each of the tract's 20 strongest connections, separately for tracts in districts with fewer than seven and more than 14 tracts. The connection strength between two tracts is defined in the same way as in the district estimation algorithm. We restrict to tract pairs that are in the same county, and we include only tracts in counties that have at least 75 tracts and are in a CBSA with population over 250,000.

Sources: Authors' calculations using district boundaries and EIF data.

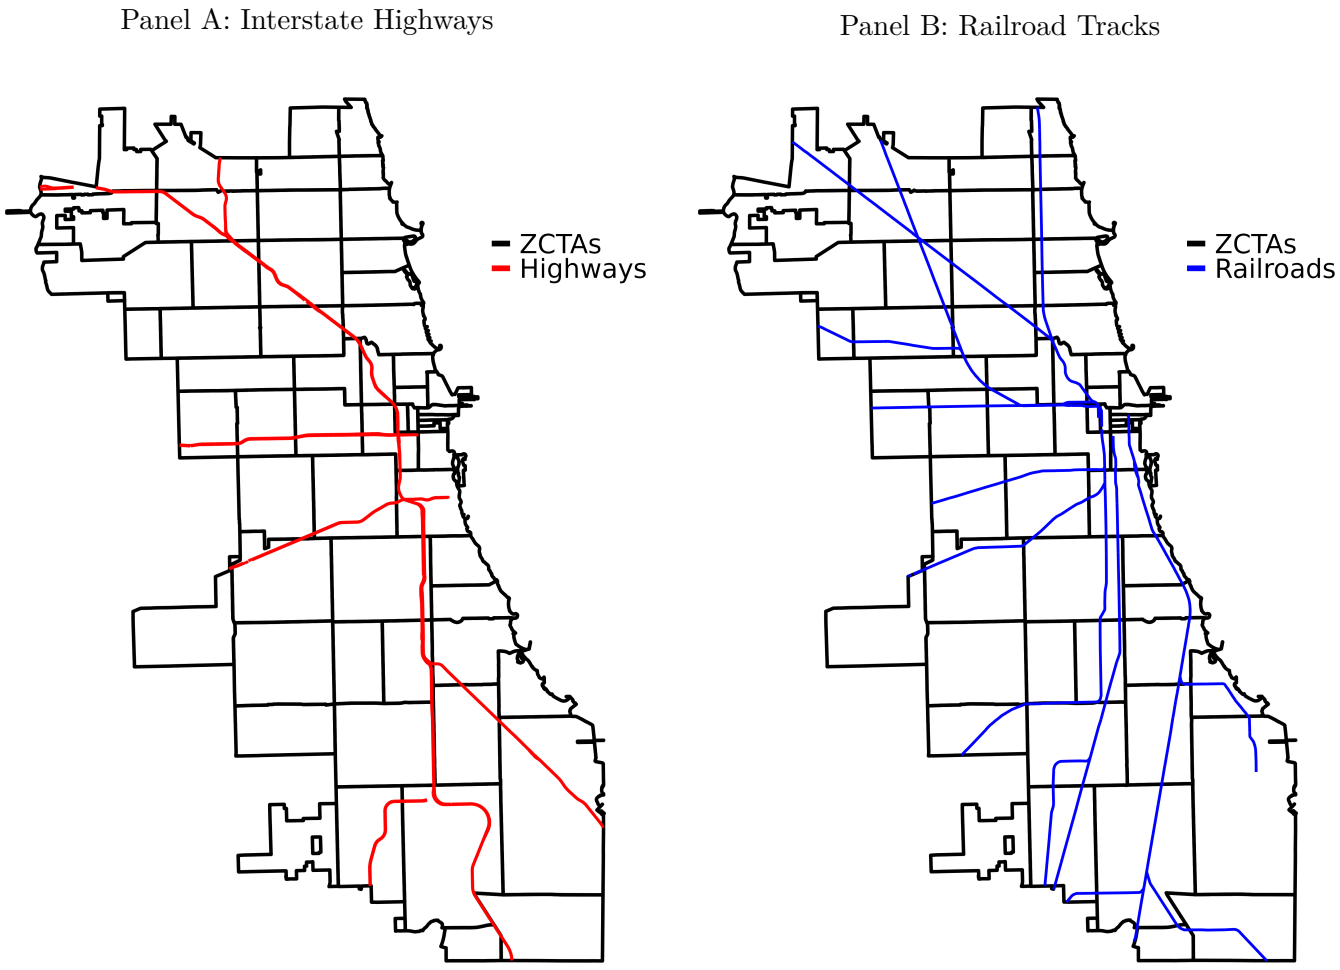
Figure A.8: Distribution of Border Frequency Across 10 Bootstrap Samples



Note: This figure shows the distribution of "border hardness" in the primary district results. We construct the figure by splitting the EIF data into ten 10 percent subsamples and estimating districts based on each subsample. For each pair of adjacent tracts in the sample CBSAs, we then calculate the number of bootstrap samples in which a border was placed between the two tracts. For example, we see that 14 percent of adjacent tract pairs have a border drawn between them in all 10 bootstrap samples. The figure is created using all CBSAs with population over 250,000 in the 2010-2014 ACS, which includes about 80 CBSAs that are not included in the main results of the paper.

Sources: Authors' calculations using the EIF, district boundaries, and U.S. Census TIGER/Line Shapefiles.

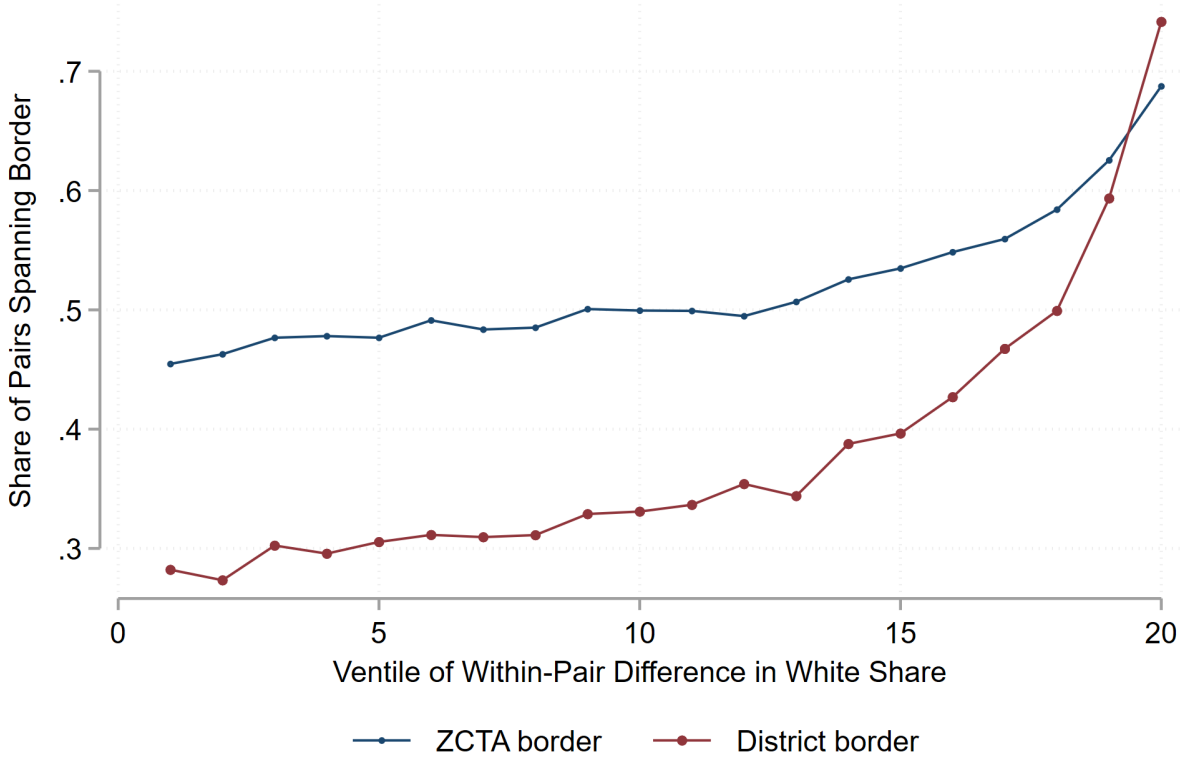
Figure A.9: Zip Code Tabulation Areas and Physical Barriers



Note: This figure plots zip code tabulation areas in Chicago against interstate highways (left panel) and major railroad lines (right panel, Chicago Transit Authority lines excluded).

Sources: City of Chicago Data Portal Railroads and Major Streets shapefiles, U.S. Census TIGER/Line Shapefiles.

Figure A.10: Share of Adjacent Tract Pairs Spanning Border Versus White Share Difference

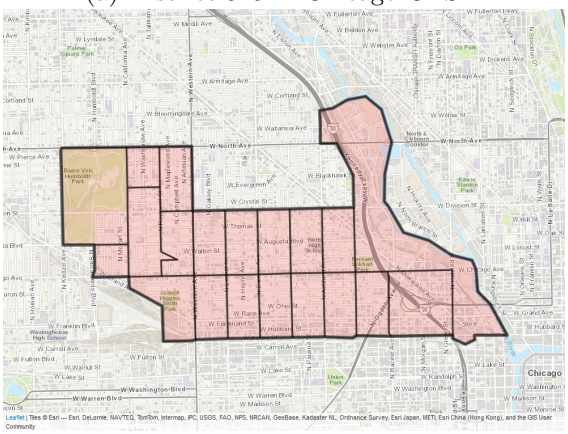


Note: This figure shows how the probability that a pair of adjacent tracts are part of different districts or ZCTAs varies with the difference in the two tracts' racial composition. The x-axis indicates the ventile of difference in the white share of the tracts in a pair, while the y-axis shows the share of pairs in each ventile that span a district and ZCTA border.

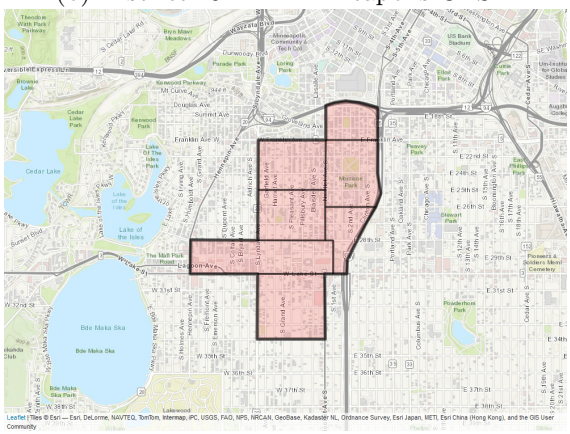
Sources: Authors' calculations using district boundaries, U.S. Census TIGER/Line Shapefiles, and demographic information compiled from the 2010 decennial census in the LTDB.

Figure A.11: Districts Containing Irregular Tracts

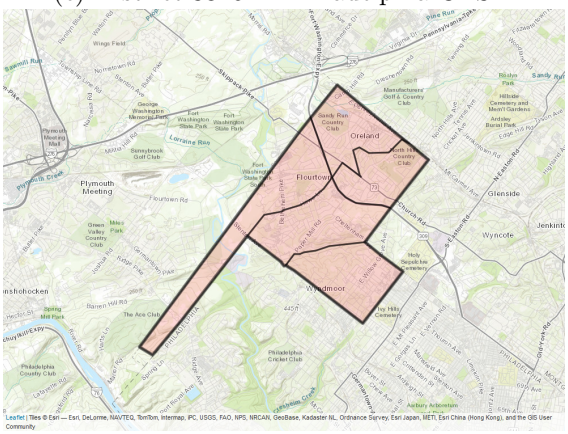
(a) District 815 in Chicago CBSA



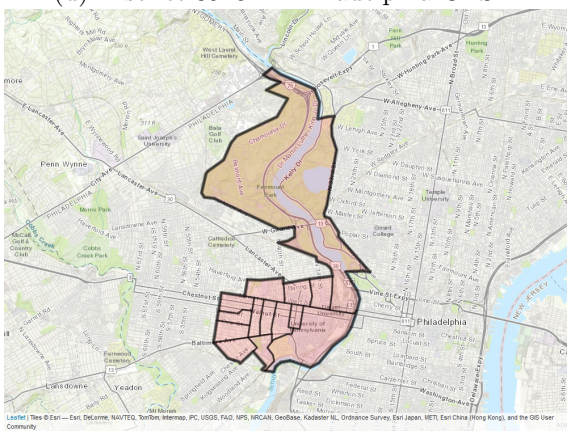
(b) District 2941 in Minneapolis CBSA



(c) District 3846 in Philadelphia CBSA



(d) District 3928 in Philadelphia CBSA



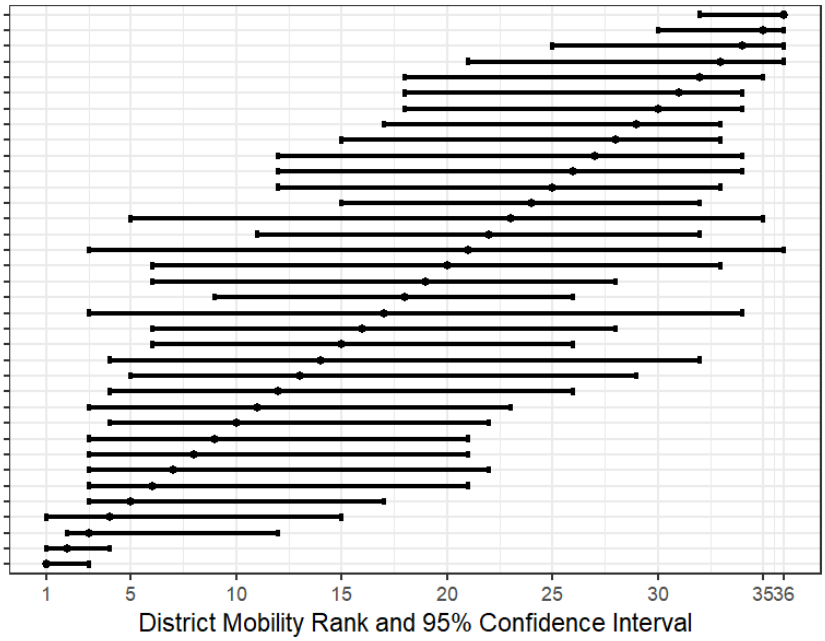
Note: This figure provides examples of several types of irregularly shaped tracts that can lead to odd district shapes. The district of interest is highlighted in pink, and the tracts that compose it are outlined in black.

In Panel A, District 815 approximates the West Town neighborhood in Chicago. The easternmost tract in the district largely consists of an interstate highway, a major rail viaduct, and an industrial area. This tract is quite large because of its low population density, and it prevents the district from matching popular perception. In Panel B, District 2941 consists of roughly the Whittier neighborhood of Minneapolis. The western boundary of the district is generally formed by Lyndale Avenue, which also forms the western boundary for a north-south column of tracts. However, one tract extends west beyond Lyndale, leading the district to have a zigzag in its western boundary.

In Panel C, District 3846 consists of exactly Springfield Township in Montgomery County, Pennsylvania. The township contains a long and narrow extension to the southeast, which was incorporated into a single similarly oddly shaped tract. In Panel D, District 3928 includes much of the University City neighborhood of Philadelphia. However, over half of its surface area consists of a single large tract that contains a small amount of housing and a large park, leading the district's overall shape to differ significantly from what locals would consider University City.

Sources: District boundaries and U.S. Census TIGER/Line Shapefiles. Basemap by Esri via Leaflet.

Figure A.12: Confidence Sets on Ranking of King County District Intergenerational Mobility



Note: This figure shows 95 percent simultaneous confidence sets on the ranking of intergenerational mobility in districts in King County, WA. Each horizontal line represents a separate district, with the diamond marker representing the rank of the district's mobility point estimate and the bar representing the confidence set on that rank. We construct district-level mobility estimates as simple aggregations of the tract estimates and their standard errors, using mobility estimates for children born to parents at the 25th income percentile. Confidence sets are constructed using the R package provided in the supplemental material of Mogstad et al. (2024).

Sources: Authors' calculations using district boundaries and Opportunity Atlas intergenerational mobility estimates.

Table A.1: Share of Moves Accounted for by Demographic Subgroups

Subgroup	*Share of Moves*
Renter	0.6435
Owner-occupant	0.3565
White	0.5003
Black	0.2039
Hispanic	0.1123
Other Race	0.1835
Age 0–18	0.3243
Age 19–29	0.2603
Age 30–65	0.3667
Age 66+	0.0486
Children in house	0.6293

Note: This table shows the share of moves from 2010 to 2011 that were accounted for by different subgroups of people. A person's subgroup is determined based on their response to the 2010 decennial census, and migration is determined using the 2010 and 2011 EIFs.

Sources: Authors' calculations using 2010 decennial census, 2010 and 2011 EIF.

Table A.2: Migration Connectivity and Distance

Panel A: Connection distribution of bordering (not bordering) tracts

	Share in connection rank bin			
	1–10	11–20	21–50	Over 51
Bordering Tract	0.8163	0.1122	0.0543	0.0172
Not Bordering Tract	0.0229	0.0372	0.1176	0.8224

Panel B: Distance distribution of strong (weak) migration connections

	Share in distance rank bin			
	1–10	11–20	21–50	Over 51
Top 10 Migration Connection	0.6042	0.2005	0.1341	0.0612
Weaker Migration Connection	0.0163	0.0325	0.1162	0.8350

Note: This table provides some information on the relationship between migratory connectivity and spatial proximity. The first row of Panel A shows the share of tracts bordering the index tract that fall into four bins of connectivity with the index tract, and the second row shows the same for tracts in the same county that do not border the index tract. The first row of Panel B shows the share of the top 10 migration connections of the index tract that fall into four bins of distance from the index tract, and the second row shows the same for the index tract's other connections. We restrict to tract pairs that are in the same county, and we include only tracts in counties that have at least 75 tracts and are in a CBSA with population over 250,000.

Sources: Author's calculations using EIF and U.S. Census TIGER/Line Shapefiles.

Table A.3: Tract Home Values and Intergenerational Mobility Versus District Characteristics

	DV: Log(Home Value)			**DV: OA Mobility**		
	(1)	(2)	(3)	(4)	(5)	(6)
𝟙(Top quartile of		-0.024			-0.011	
[distr pov - tract pov])		(0.004)			(0.001)	
Tract share Black	-0.121	-0.120	0.007	-0.112	-0.112	-0.091
	(0.023)	(0.023)	(0.031)	(0.003)	(0.003)	(0.004)
Tract share white	0.117	0.108	0.187	0.013	0.009	0.034
	(0.022)	(0.022)	(0.025)	(0.003)	(0.003)	(0.004)
Log tract income	-1.339	-1.324	-1.947	-0.127	-0.119	-0.130
	(0.212)	(0.212)	(0.190)	(0.029)	(0.029)	(0.030)
(Log tract income)^2	0.076	0.075	0.105	0.007	0.006	0.007
	(0.010)	(0.010)	(0.009)	(0.001)	(0.001)	(0.001)
Tract college share	1.129	1.129	0.922	0.051	0.051	0.066
	(0.023)	(0.023)	(0.025)	(0.004)	(0.004)	(0.004)
Tract unemp. share	-0.540	-0.538	-0.279	-0.012	-0.012	-0.019
	(0.054)	(0.054)	(0.045)	(0.006)	(0.006)	(0.005)
R^2	0.846	0.846	0.898	0.667	0.670	0.772
Observations	45,600	45,600	45,600	45,600	45,600	45,600
Tract poverty spline	Y	Y	Y	Y	Y	Y
County fixed effects	Y	Y	Y	Y	Y	Y
District fixed effects	N	N	Y	N	N	Y

Note: This table shows results from regressing the log of tract median home value (left three columns) or expected adult income of children of parents at the 25th percentile of income (right three columns) on different combinations of tract and district characteristics. Each column represents a separate regression. Columns one and four provide a baseline that uses only tract-level information. In columns two and five, we add an indicator for whether a tract is in the top quartile of the difference between the tract's poverty rate and the mean poverty rate among the other tracts in its district. Columns three and six instead show much variation can be explained by adding district fixed effects. The tract poverty spline consists of bins 5 percentage points wide, top-coded at 40 percent.

Sources: Authors' calculations using district boundaries, intergenerational mobility estimates from Chetty et al. (2025), and demographic information compiled from the 2010 decennial census (racial composition) and 2008-2012 American Community Survey aggregates (all other variables) in the LTDB.

Table A.4: Specification Search Results

Connection weight	Exponent	Cutoff	Mean tract count	Share correct size	Share regularized	Share discontiguous	Bootstrap stability	Share disconnected (midsize)	Share disconnected (large)
Move count	1	10	72.3	0.276	0.002	0.363	0.722		
Move count	2	10	13.2	0.723	0.021	0.156	0.806		
Move count	3	10	9.7	0.774	0.044	0.114	0.818		
Move count	4	10	8.3	0.773	0.064	0.112	0.808		
Inverse rank	1	10	12.6	0.872	0.007	0.116	0.850	0.201	0.297
Inverse rank	2	10	9.4	0.909	0.018	0.088	0.884	0.166	0.232
Inverse rank	3	10	7.3	0.858	0.050	0.084	0.848	0.199	0.397
Inverse rank	4	10	6.2	0.770	0.099	0.098	0.799	0.292	0.548
Inverse rank	1	5	8.6	0.919	0.020	0.087	0.859	0.123	0.313
Inverse rank	2	15	15.9	0.731	0.004	0.141	0.876	0.182	0.248

Note: This table shows the performance of several Infomap parameterizations. Share correct size shows how many districts contain between five and 20 tracts. Share regularized and share discontiguous represent the share of tracts that initially either fell below the minimum size threshold or were discontiguous, respectively. The bootstrap stability column shows the share of adjacent tract pairs that have the same district border status in an estimation using the full sample and in an estimation using a random 10 percent sample. In the share disconnected column, we label a district as disconnected if it contains distinct clusters of tracts that are not strongly connected to each other by migration. We report this share separately for midsize districts (8-10 tracts) and large districts (15-20 tracts). Details on how we construct this connectivity measure are included in Section A.1 of the online appendix. The table is created using districts in all CBSAs with population over 250,000, which includes about 80 CBSAs that are not included in the main results of the paper.

Sources: Authors' calculations using EIF, population from 2010-2014 ACS.

Table A.5: Border Agreement Across Specifications

Bootstrap Border Count	Inv. Rank X=1, C=10	2000s Moves	Public Release	Inv. Rank X=2, C=15	Inv. Rank X=1, C=15	Inv. Rank X=3, C=5
0	0.9883	0.9549	0.9755	0.9864	0.9673	0.9326
1	0.9020	0.8306	0.9193	0.8917	0.8694	0.8078
2	0.8055	0.7075	0.8619	0.7742	0.8061	0.7597
3	0.7014	0.6505	0.8177	0.6505	0.7645	0.7404
4	0.6145	0.5756	0.7938	0.5360	0.7471	0.7462
5	0.5590	0.5918	0.8077	0.4369	0.7430	0.7740
6	0.5768	0.6016	0.8257	0.4510	0.7411	0.7998
7	0.6224	0.6798	0.8469	0.4799	0.8139	0.8429
8	0.7165	0.7427	0.8936	0.5321	0.8701	0.8791
9	0.7977	0.8351	0.9517	0.6315	0.9363	0.9280
10	0.9309	0.9515	0.9893	0.8093	0.9885	0.9841

Note: This table shows, for a number of alternative specifications, the share of adjacent tract pairs that have the same district border status in the primary specification and the alternative indicated in the column heading. The rows report this agreement within levels of "border hardness," defined as the number of bootstrap samples in which a border was placed between the two tracts in the pair (under the primary specification). For example, the top left cell shows that among tract pairs that were never assigned a district border in any bootstrap samples, 98.8 percent have the same district border status in the baseline primary specification and in an estimation where the exponent X is set to one and the cutoff C is set to 10. See Section A.1 of the online appendix for details on the alternative model specifications referenced in each column. The 2000s Moves column refers to districts estimated on moves between 2000 and 2009, while the Public Release column refers to the districts estimated using noise-infused migration flows (the version that we publicly release). The table is created using districts in all CBSAs with population over 250,000, which includes about 80 CBSAs that are not included in the main results of the paper.

Sources: Authors' calculations using EIF, U.S. Census TIGER/Line Shapefiles, population from 2010-2014 ACS.