# Heterogeneous Preferences for Neighborhood Amenities: Evidence from GPS Data<sup>\*</sup>

 $Cody Cook^{\dagger}$ 

August 2024

#### Abstract

This paper examines how preferences for neighborhood amenities vary by income. Using panel data on over 100 million visits to 1.4 million establishments, I build and estimate a discrete choice model of demand for restaurants, shops, personal services, and entertainment places. While preferences for specific establishments often vary by income, preferences for each neighborhood's overall access to amenities are highly aligned. Dense urban areas have a sufficient variety of amenities to offer broad appeal, while less dense areas have more limited access to amenities. For incumbent residents of gentrifying neighborhoods, counterfactual simulations suggest that the tailoring of amenities to higher-income entrants has only modest welfare effects relative to the effects of displacement to cheaper neighborhoods with worse access to amenities.

<sup>\*</sup>I am grateful for valuable feedback from Lanier Benkard, Victor Couture, José Ignacio Cuesta, Lindsey Currier, Rebecca Diamond, Liran Einav, Matthew Gentzkow, Ed Glaeser, Jessie Handbury, Pearl Li, Neale Mahoney, Peter Reiss, Brad Ross, Paulo Somaini, Winnie van Dijk, Shoshana Vasserman, and many seminar participants. This project began while I was a contractor for the urban data platform company Replica, which provided the data and necessary computational resources; Replica did not ask for nor receive any editorial oversight in either the topic selection or the analyses conducted. Steven Kim, Dave Lawlor, Kiran Jain, Aude Marzuoli, and Alexei Pozdnoukhov at Replica all provided valuable assistance.

<sup>&</sup>lt;sup>†</sup>Yale University. Email: codyfcook@gmail.com

# 1 Introduction

Urban neighborhoods have undergone significant transformations in recent decades, with higher-income residents increasingly moving to previously low-income areas (Baum-Snow and Hartley, 2020; Couture and Handbury, 2020). This trend—often connected to 'gentrification'<sup>1</sup>—has sparked debates over its impacts on long-time residents and urban inequality (Ellen and Ding, 2016). The role of amenities, such as restaurants and shops, often features prominently in both the academic and popular discourse on gentrification.<sup>2</sup> If new amenities in a gentrifying area cater solely to the preferences of higher-income entrants, they may amplify spatial sorting and create 'Tiebout clusters' of neighborhoods for the rich and others for the poor (Tiebout, 1956; Bayer and McMillan, 2012; Almagro and Dominguez-Iino, 2024). Conversely, if the new amenities have broad appeal, incumbents displaced by rising rents may miss out on improvements to their neighborhood's amenities (Couture et al., 2023).

In this paper, I study how preferences for neighborhood amenities vary by income. I build a discrete choice model of demand for restaurants, shops, personal services, and entertainment places. I estimate the model using data on over 100 million visits to 1.4 million unique establishments amenities, which I infer from the location histories of 7.7 million GPS devices in the 30 largest MSAs.<sup>3</sup> In contrast to approaches that use residential choice data to estimate the aggregate value of a neighborhood's amenities,<sup>4</sup> data on individual visits allow me to estimate revealed preferences down to the level of a specific establishment. Using the estimated model, I evaluate the correlation structure of preferences for both establishments and neighborhoods, and the extent to which observable characteristics can explain

<sup>&</sup>lt;sup>1</sup>The word 'gentrification' was coined by Ruth Glass in 1964 to describe the influx of upper-class households (the 'gentry') that displaced working-class residents of the urban core in London (Glass, 1964)

<sup>&</sup>lt;sup>2</sup>See, for example, recent articles in the New York Times (Gordinier, 2016; Kolomatsky, 2020), the Wall Street Journal (Raleigh, 2017; Ukueberuwa, 2020), and the Atlantic (Smith, 2016)

<sup>&</sup>lt;sup>3</sup>Metropolitan Statistical Areas (MSAs) are defined by the US Census Bureau and generally consist of an urban core and surrounding counties that have close economic ties to the core. I select the 30 largest MSAs based on their 2010 populations.

<sup>&</sup>lt;sup>4</sup>This approach was introduced by Rosen (1981) and Roback (1982), and remains common today in residential choice models such as Bayer and Timmins (2007) and in the quantitative spatial economics literature, reviewed in Redding and Rossi-Hansberg (2017).

preference differences. I then assess the potential welfare impact of two common aspects of gentrification: the tailoring of amenities to higher-income residents and the displacement of long-time residents to cheaper neighborhoods.

I begin by documenting descriptive patterns of amenity access and consumption by income. To measure income, I match each device in the GPS data to its home parcel and predict income based on the parcel characteristics (e.g., market value, size, and location). Residents in the top quartile of estimated household income visit 30% more restaurants, 15% more personal services, and 45% more shops and entertainment places each quarter than residents in the bottom quartile, despite having fewer amenities of each category within a 10-minute drive from home. Even for residents of the same block group, those in the top income quartile visit 8-22% more amenities of each category than those in the bottom quartile. Higher-income residents also visit different types of establishments within a category. For example, visitors to chains such as Costco, Chipotle, Whole Foods, and CrossFit have higher estimated incomes than visitors to chains such as Dollar General, Taco Bell, Save-A-Lot, and Planet Fitness. These descriptive differences, however, conflate preferences and access – while consumption of amenities varies by income, so does access.

To disentangle preferences from access, I model demand for amenities as a nested logit, in which residents first select a category of amenities (e.g., shops), then a subcategory (e.g., clothing stores), and finally a specific establishment. I partition establishments into 4 main categories—restaurants, shops, personal services, and entertainment—and 33 subcategories. The data do not include many demand-relevant characteristics, so I specify the model in 'product-space' and estimate establishment-level fixed effects. I also do not observe prices, and some amenities, such as parks, are free. As a measure of cost, I use the driving time from an individual's home and work and convert all measures of utility into units equivalent to saving a minute of driving time. Intuitively, establishments for which residents either visit more often or travel farther to visit must be higher value. I allow preferences to vary by whether each device's estimated household income is above ('higher-income') or below ('lower-income') the median household income in the sample MSAs.

Higher and lower-income residents' preferences for establishments are positively correlated in all subcategories, but the degree of correlation varies. For large retail subcategories such as malls and general merchandise stores, preferences have a correlation coefficient of over 0.75. Within these subcategories, most establishments provide similar value across the income distribution. Preferences for subcategories such as restaurants, barber shops, and gyms are less correlated, and establishments are more likely to appeal to only part of the income distribution. Observable characteristics of an establishment—even just its subcategory—are often informative about its appeal to residents. Many entertainment subcategories provide more overall value to higher-income residents, while many retail subcategories provide more value to lower-income residents, perhaps because they are less likely to shop online (Dolfen et al., 2023). For a subset of establishments, I observe more detailed characteristics from their Yelp page and find that, for example, higher-income residents disproportionately prefer restaurants with higher Yelp price levels and cuisines like New American.

At the neighborhood level, however, differences between income groups tend to wash out. Instead, neighborhoods are generally either sufficiently dense in amenities that they offer broad appeal or have a more limited set of amenities to offer. To show this, I construct a Neighborhood Amenity Quality Index (NAQI), which measures the value of amenity access for higher and lower-income residents of different Census block groups using a variant of the log-sum measure of consumer surplus for discrete choice models (Small and Rosen, 1981).<sup>5</sup> For the average household, moving from the 25th to 75th percentile of the NAQI distribution is worth the equivalent of saving 85 minutes of driving per week. Dense neighborhoods especially those near the urban core—benefit from residential agglomeration effects and offer especially valuable access to amenities. Each doubling of population density is associated with an increase in value of almost 40 minutes/week, of which about 45% is due to saved travel times. The strong relationship between density and the NAQI values echoes results

<sup>&</sup>lt;sup>5</sup>The NAQI data are available to download here.

from Ahlfeldt et al. (2015), who find that agglomeration forces are twice as strong for residential amenities than for productivity. Conditional on density and proximity to the core, neighborhoods with higher household income, rent, and education tend to have more valuable access to amenities, especially for higher-income households.

The degree of alignment in higher and lower-income preferences over neighborhoods suggests that displacement from a high-amenity neighborhood may have a larger welfare impact than the tailoring of neighborhood amenities to a specific income group. To evaluate this channel, I use individual-level migration data from Infutor to identify long-time residents of high-amenity neighborhoods who move to a cheaper neighborhood after their own neighborhood experiences a large rent increase. On average, these residents move to a neighborhood with 30% lower median rents, but with a NAQI value that is 39 minutes/week lower. In contrast, tailoring amenities to higher-income residents has only modest effects. I simulate a counterfactual world in which the top 25% of establishments for lower-income residents in each subcategory are replaced by replicas of the top 25% of establishments for higher-income residents. Even under this fairly extreme version of tailoring, the average NAQI value increases by just 8 minutes/week for higher-income residents and decreases by 7 minutes/week for lower-income residents. These findings suggest that policies aimed at preventing residential displacement may be more beneficial for incumbents than those focused on maintaining specific types of neighborhood amenities.<sup>6</sup>

**Related literature.** This paper contributes to three main strands of literature. The first studies the causes and consequences of gentrification and neighborhood change, often focusing on role of neighborhood amenities and the "consumer city" (Glaeser, Kolko and Saiz, 2001).<sup>7</sup> Recent work by Couture and Handbury (2020) and Baum-Snow and Hartley (2020)

<sup>&</sup>lt;sup>6</sup>Examples of policies that target the composition of a neighborhood's amenities include zoning restrictions, bans on chain stores (Klopack, 2024), and the designation of landmarks or cultural heritage districts. <sup>7</sup>For earlier work on gentrification more broadly, see, for example, Vigdor (2002); Ellen and O'Regan (2010); McKinnish, Walsh and Kirk White (2010); Guerrieri, Hartley and Hurst (2013); Lester and Hartley (2014); Meltzer and Ghorbani (2017); Brummet and Reed (2021); Berkes and Gaetani (2023); Oh and Seo (2023). Much of this literature is reviewed in Couture and Handbury (2023).

shows that the return of higher-income, college-educated residents to urban cores is due to a rising tendency for these residents to sort towards areas rich in consumption amenities. Such trends may reduce the welfare of lower-income incumbents if the amenities tailor to the preferences of the entrants (Waldfogel, 2003, 2010; Almagro and Dominguez-Iino, 2024), or if incumbents are displaced by rising rents (Pennington, 2021; Couture et al., 2023). I find that the latter channel is likely of greater importance: lower-income residents still value the amenity access of the urban core more than that of other neighborhoods. The verticality of preferences for neighborhoods also helps validate the common modeling assumption of a single index of neighborhood amenities over which preferences may be stronger for higherincome residents (see, e.g., Couture et al., 2023).

Second, this paper contributes to methodological work on estimating the value of various neighborhood amenities and how it may vary by individual characteristics. A common approach is to infer value based on residential choices. Using this approach, researchers have shown that higher-income households place more value on school quality (Bayer and Timmins, 2007), other high-skill neighbors (Diamond, 2016), reduced crime (Ellen and O'Regan, 2010), proximity to work (Su, 2022; Barwick et al., 2024), and distance from public housing (Almagro, Chyn and Stuart, 2023). Instead, I take a 'bottom-up' approach and use consumption data to estimate revealed preferences for specific establishments.<sup>8</sup> This latter approach builds on work by Couture (2016), who shows that the returns to living in areas dense in restaurants are more due to increased variety than to saved travel time, and Davis et al. (2019), who show how both spatial and social frictions lead to racial segregation in restaurant consumption. My primary contribution to this literature is one of scale: by combining large-scale GPS data with computational tools borrowed from the machine learning world, I estimate preferences for a range of amenity categories down to the level of specific establishments. The estimates highlight, for example, how certain types of amenities with

<sup>&</sup>lt;sup>8</sup>Similar work by Handbury and Weinstein (2015) and Handbury (2021) use data on product-level choices to estimate local price indices. Caetano and Maheshri (2019) use data from Foursquare to show that men and women visit different types of establishments.

less aligned preferences (e.g., personal services) have more potential to act as amplifiers of spatial inequality than other types with more aligned preferences (e.g., general merchandise).

Finally, this project builds on a growing literature that uses data from GPS-enabled devices to study questions related to urban mobility. Applications of GPS data include mobility during the COVID-19 pandemic (Chang et al., 2021; Allcott et al., 2020), waiting times at voting polls (Chen et al., 2022), the effects of new transit options (Gupta, Van Nieuwerburgh and Kontokosta, 2022; Cook and Li, 2023), alternative measures of segregation based on where people spend their time (Caetano and Maheshri, 2019; Athey et al., 2021), and differences in urban mobility by students and adults (Cook, Currier and Glaeser, 2024). Most related to my work, Athey et al. (2018) estimate demand for lunch restaurants and Miyauchi, Nakajima and Redding (2022) study how the consumption externalities of commute patterns shape the spatial structure of cities. A common downside of GPS data is the lack of individual-level demographics. In this project, I use a novel method for inferring demographics using information on each device's home parcel.

# 2 Data

I identify amenities using SafeGraph's Places dataset, which includes the name, category, geographic footprint, and other characteristics of many Points of Interest (POIs). The data encompass a wide range of both commercial and public amenities such as restaurants, shops, barbers, gyms, parks, churches, zoos, libraries, dentists, and banks, but excludes many other features of a neighborhood that residents may appreciate, such as trees, clean streets, and access to quality schools.<sup>9</sup> I use each establishment's North American Industry Classification System (NAICS) code to categorize them into 4 main categories—restaurants, shops, personal services, and entertainment places—and 33 subcategories.<sup>10</sup> For example, shops are

<sup>&</sup>lt;sup>9</sup>I use 'establishment' to refer to the location of any neighborhood amenity, including both private establishments and public amenities such as parks. In Appendix Section A.3.2, I show that the SafeGraph data has similar coverage to the County Business Patterns from the Census.

<sup>&</sup>lt;sup>10</sup>Appendix Table A.1 documents the NAICS codes used to identify each subcategory and the corresponding number of establishments.

divided into 13 subcategories, including grocery stores, malls, and clothing stores. Where available, I augment the data using characteristics scraped from Yelp, such as price levels. The final sample covers 1.4 million POIs.

Next, I construct a panel of visits to amenities during 2019 by matching data on the location histories of GPS-enabled devices to the footprints of each POI. The raw GPS data consist of 'pings' with time and location information, which can be triggered either by opening a specific app or via apps that share location in the background. The pings are then combined into 'stays' at different coordinates. For example, a visit to the supermarket may result in a hundred pings, all of which would become a single stay. Using these stays, I assign home and work locations to each device-quarter using heuristics on when people are at home and work and exclude visits to establishments at the same location as a device's home or work. For each establishment, I compute the driving time from devices' home and work locations using Graphhopper, a router built on OpenStreetMaps.<sup>11</sup> I describe the data processing steps in greater detail in Appendix A.

GPS data have both advantages and disadvantages relative to other sources of individual consumption patterns, such as credit card and debit card data (Dolfen et al., 2023; Klopack, 2024; Relihan, 2024). On the advantages side, GPS data include visits to amenities that either do not involve a transaction (e.g., a public park) or for which the individual paid in cash. According to a survey by Foster, Greene and Stavins (2020), 30% of in-person retail transactions are paid for in cash, with lower-income households being disproportionately likely to pay in cash. GPS data providers also put fewer restrictions on the use of the data, so researchers can disclose statistics about specific brands, e.g., the proportion of visits to Whole Foods versus Safeway.

The primary disadvantage of GPS data is that they do not include any demographic information. A common solution is to infer demographics based on the median household

<sup>&</sup>lt;sup>11</sup>To reduce the number of routing requests, I compute driving times between pairs of block groups within each MSA (over 375 million total routes) rather than door-to-door. Graphhopper returns only the free-flow driving times, which do not account for traffic.

in a device's home Census tract. However, using median income introduces substantial measurement error for devices living in mixed-income tracts. To improve precision, I combine information on a device's home parcel from CoreLogic (e.g., market value, bedroom size, and building type) and the income distribution of the surrounding block group to estimate household income for each device. I describe this procedure in more depth in Appendix Section A.2 and discuss robustness to alternative measures of income in Section 5.4.

Past researchers have found that GPS data is broadly representative (see, e.g., Couture et al., 2022), although it tends to slightly over-sample from less dense, younger, and less white block groups within a city. I discuss several measures of sample quality in Appendix Section A.3. For all analyses, I use sample weights to adjust for non-uniform sampling across space. I weight each device-quarter by  $\omega_{iq} = N_{g(iq)}^{Total}/N_{g(iq)}^{GPS}$  where  $N_{g(iq)}^{GPS}$  is the number of GPS devices observed in the device's home block group g(iq) and  $N_{g(iq)}^{Total}$  is the 2019 population of the block group in the American Community Survey (ACS) 5-year tabulations.

Table 1 summarizes the sample used for analyses. The final sample consists of 7.7 million devices and 1.2 billion stays. For the average device-quarter, I observe 95 stays, of which 48 are at home, 12 are at work, and 12 are at one of the neighborhood amenities in the sample.

# 3 Access and consumption

This section presents three empirical facts about the relationship between income and amenity access and consumption.

#### Fact 1: access to amenities is generally decreasing in income

Higher-income residents often live—and, to a lesser extent, work—near fewer amenities of each category. For each individual in the data, I compute the number of establishments in each category within a 10-minute driving radius from their home and work locations. I then regress the log number of establishments on indicators for quartile of estimated household income and MSA fixed effects. On average, individuals in the bottom quartile of household income have 25-40% more shops, restaurants, and personal services within a 10-minute radius than individuals in the top half of the income distribution (Figure 1 Panel a). Differences in access from work are more muted, with the top and bottom quartiles of income having similar access and the middle two quartiles having 5-15% fewer establishments of each category.

Population density is a key determinant of access to amenities from home. While many higher-income residents have moved towards city centers in recent decades, most continue to live in less dense suburbs with worse access to amenities (Couture et al., 2023). When comparing individuals in similarly dense areas by adding controls for log population density, the relationships between income and access to amenities of different categories are diminished by over 70%, and even reversed for entertainment establishments (Appendix Figure C.1).

#### Fact 2: consumption of amenities is increasing in income for all categories

Despite often having worse access, higher-income residents consume more amenities in their city. To show this, I regress the log number of visits to amenities of each category on indicators for quartile of estimated household income and MSA fixed effects. Compared to those in the bottom quartile of household income, individuals in the top quartile make 30% more visits to restaurants, 46% more to shops, 16% more to personal services, and 47% more to entertainment places (Figure 1 Panel b). Even with additional fixed effects for an individual's home and work block group, individuals in the top income quartile continue to make 8-22% more visits per quarter to each amenity category than those in the bottom quartile. The differences in consumption behavior within a given block group help validate using parcel-level rather than block group-level income estimates.

Amenity consumption varies by demographics beyond income. In Appendix Figure C.2, I investigate differences by race and education. For race, I follow Athey et al. (2021) and classify each device as a 'white device' (WD) or 'non-white device' (NWD) based on whether their home block is more or less than 50% non-Hispanic white in the 2010 Census. For education, I estimate whether a device is likely to be college-educated using the characteristics of their home parcel. Within each demographic group, similar differences emerge across income quartiles, with the top income quartile consuming 15-50% more of each category than the bottom quartile. Consumption levels, however, differ. WDs consume more of each amenity category than NWDs at all income levels, and college-educated devices consume more restaurants and entertainment amenities at all income levels than non-college devices. While I focus on income for the remainder of the paper, each income group is best thought of as a bundle of characteristics that vary by income, such as race, education, and age.<sup>12</sup>

#### Fact 3: consumption varies by observables, such as chain affiliation

Certain chains, such as Whole Foods and Starbucks, are commonly associated with gentrification in the popular press (see, e.g., Smith, 2016; Kolomatsky, 2020). Figure 2 documents the average visitor income, visitors' distance from home, and the number of visits per establishment for popular brands in four subcategories dominated by large chains. For general merchandise, visitors to Costco, BJ's, and Sam's Club have higher estimated household income than visitors to Walmart and many dollar store chains.<sup>13</sup> For fast food, visitors to Five Guys and Chipotle have higher incomes than visitors to Burger King and Taco Bell. For grocery stores, Whole Foods, Trader Joe's, and Safeway all attract especially high-income consumers. Finally, for fitness centers—where the average income at any establishment is generally higher than for other subcategories—visitors to premium gyms such as CrossFit, Equinox, and Orangetheory have average household incomes over \$115,000.

Chains also vary in the number of visitors they attract and how far the average visitor travels from home. Popular chains such as Costco, Chick-fil-A, Lifetime Fitness, and Whole Foods attract more visitors and visitors from farther away than the average chain in the same subcategory. As I turn to the model in the next section, both sources of variation will be key for identifying the value of specific establishments.

 $<sup>^{12}\</sup>mathrm{Among}$  devices with above-median estimated household income, 54% are estimated to have a college degree and 75% are WDs. For devices with below-median income, only 5% are estimated to have a college degree and 44% are WDs.

 $<sup>^{13}</sup>$ Cao (2022) shows that dollar stores are often located closer to lower-income residents. Contemporaneous work by Cao et al. (2024) finds similar differences in the average income of clientele at general merchandise chains, also using GPS data.

# 4 Model and estimation

The descriptive facts highlight differences in amenity access and consumption by income. To separate preferences for amenities from access to them, I develop and estimate a demand model. The model provides a framework for studying correlations in preferences, the value of each neighborhood's access to amenities, and the welfare impacts of changing the spatial distribution of either people or amenities.

## 4.1 A model of amenity choice

I model demand for amenities as a three-level nested logit. An individual first chooses between visiting a restaurant, shop, personal service, entertainment place, or her outside option. Second, she chooses a subcategory within the chosen category (e.g., grocery stores). Finally, she chooses a specific establishment.

**Utility.** Individuals are characterized by their income group (k) and home and work locations. The indirect utility for individual i of group k for visiting establishment j at time of week t is given by:

$$u_{ijt}^{k} = \alpha_{jt}^{k} - \kappa_{t}^{k} \mathbf{d}_{ij} + \varepsilon_{ijt}$$

$$\tag{4.1}$$

where  $\alpha_{jt}^k$  is an establishment-time fixed effect,  $\mathbf{d}_{ij}$  is vector containing the driving time from home and work, and  $\varepsilon_{ijt}$  is an idiosyncratic shock drawn from a Generalized Extreme Value (GEV) distribution that matches the nesting structure. Preferences vary by the time of week, which I discretize into bins of (weekday, weekend) × (morning, afternoon, evening).<sup>14</sup> I parameterize the  $\alpha_{jt}^k$  into the sum of fixed effects for the category  $(m_j \in \mathcal{M})$ , subcategory  $(l_j \in \mathcal{L})$ , and establishment (j):

$$\alpha_{jt}^k = \gamma_{m_jt}^k + \gamma_{l_j}^k + \gamma_j^k \tag{4.2}$$

<sup>&</sup>lt;sup>14</sup>For times of week, morning is 2am to 11am, afternoon is 11am to 5pm, and evening is 5pm to 2am the next day. I consider Friday evenings as weekend evenings and Sunday evenings as weekday evenings.

where I allow preferences for establishments to vary by the time of week only at the category level (e.g., consumers may value restaurants more in the evenings).

I formulate utility in 'product space'—i.e. with fixed effects for each establishment rather than 'characteristics space' because the data do not include many demand-relevant observables. I often only observe an establishment's name, location, and subcategory. I also do not observe prices, and use the driving time from home and, where applicable, work as the cost of consuming at a given establishment. The  $\alpha_{jt}^k$  capture the 'price inclusive' value of establishment j to an average user of group k. If an establishment is frequently visited by residents who live/work far away, it must have a high  $\alpha_{jt}^k$  term to rationalize this behavior.

I normalize the scale of utility by fixing the coefficient on driving a minute from home on weekday evenings to 1.<sup>15</sup> I also allow the variance of the idiosyncratic component ( $\varepsilon_{ijt}$ ) to differ across subcategories, such that idiosyncrasies in individual preferences can play a larger role in some subcategories than in others. With the normalization, the difference between the  $\gamma_j^k$  of two establishments corresponds to the number of minutes an average individual of group k would be willing to drive to go to one establishment over another.

A key assumption embedded in this formulation is that each visit to an amenity is decided independently. Treating each visit as independent will bias upward the estimated value of establishments frequently visited as part of a chain of stops (e.g., an ice cream shop next to a group of retail stores). An alternative approach would be to model the choice of an entire sequence of visits during a day, allowing the driving time to be amortized over multiple stops. However, such an approach introduces substantial complexity and would make it infeasible to estimate establishment-level preferences for a large choice set.<sup>16</sup> I discuss challenges with

<sup>&</sup>lt;sup>15</sup>While I use a minute of driving as a numeraire, not all trips are by car, and car usage may vary by income and destination. As such, the driving times are best interpreted as an approximation for travel times, similar to the commonly-used crow-flies distance. Recall, also, that the driving times from Graphhopper are under free-flow traffic conditions. While some of the differences in traffic across time will be captured in the time-varying coefficients, traffic also varies across space, and this approach may overstate the value of origins with above-average congestion.

<sup>&</sup>lt;sup>16</sup>Contemporaneous work by Miyauchi, Nakajima and Redding (2022) and Relihan (2024) model individual's sequence of choices in a day, but aggregate the choice set to broad categories (e.g., all non-tradeables in a neighborhood) rather than considering visits to specific establishments.

identifying 'trip chains' more in Appendix Section A.4, and conduct a robustness test using just trips that start and end at home in Section 5.4. While about 20% of observed trips include stops at multiple establishments, the average number of stops is similar for higher and lower-income residents, so it is unlikely to affect results differentially.

Choice probabilities. The probability individual i of group k chooses establishment j can be decomposed into three probabilities, one for each nested logit level (Train, 2009):

$$P_{ijt}^{k} = P_{itm_{j}}^{k} \times P_{itl_{j}|m_{j}}^{k} \times P_{itj|l_{j}}^{k}$$

$$\tag{4.3}$$

where  $P_{itm_j}^k$  is the probability of choosing category  $m_j$ ,  $P_{itl_j|m_j}$  is the conditional probability of choosing subcategory  $l_j$ , and  $P_{itj|l_j}^k$  is the conditional probability of choosing establishment j. Under the GEV assumption for  $\varepsilon$ , each of these probabilities has a closed-form solution.

Starting with the lowest level—the choice of an establishment within a subcategory—the probability follows the familiar logit form from McFadden (1973):

$$P_{itj|l_j}^k = \frac{\exp(\gamma_j^k - \boldsymbol{\kappa}_t^k \mathbf{d}_{ij})}{\sum\limits_{j' \in l_j} \exp(\gamma_{j'}^k - \boldsymbol{\kappa}_t^k \mathbf{d}_{ij'})}$$
(4.4)

which depends only on each establishment's location relative to the individual's home and work locations and the component of  $\alpha_{jt}^k$  that affects choices within a subcategory  $(\gamma_j^k)$ .

In the upper levels of the model, an individual's choice of a subcategory depends on her access to establishments of each subcategory and the value of those establishments for her income group. The probability of consuming at *any* establishment within subcategory l is given by

$$P_{itl}^{k} = P_{itm_{l}}^{k} * P_{itl|m_{l}}^{k}$$

$$= \underbrace{\left(\frac{\exp\left[\gamma_{m_{l}t}^{k} + \rho_{m_{l}}^{k} \mathrm{IV}_{itm_{l}}^{k}\right]}{1 + \sum_{m' \in \mathcal{M}} \exp\left[\gamma_{m't}^{k} + \rho_{m'}^{k} \mathrm{IV}_{itm'}^{k}\right]}\right)}_{\text{Choice of category}} \underbrace{\left(\frac{\exp\left[(\gamma_{l}^{k} + \beta^{k} \mathrm{DTE}_{itl}^{k})/\rho_{m_{l}}^{k}\right]}{\sum_{l' \in m_{l}} \exp\left[(\gamma_{l'}^{k} + \beta^{k} \mathrm{DTE}_{itl'}^{k})/\rho_{m_{l'}}^{k}\right]}\right)}_{\text{Choice of category}}$$

$$(4.5)$$

where  $m_l$  denotes the category containing subcategory l,  $\rho_{m_l}^k$  measures the degree of independence in unobserved utility for subcategories within a category, and the utility of the outside option is normalized to zero. Two sets of inclusive values link the three levels of the nested logit: the inclusive value of subcategories within each category ( $IV_{itm}^k$ ) and of establishments within each subcategory ( $DTE_{itl}^k$ ), which, given the normalization using the disutility of driving from home, I refer to as 'driving time equivalents,' or DTEs.<sup>17</sup>

## 4.2 Estimation

The model makes several assumptions and simplifications in the pursuit of tractability at scale. Even so, the number of observed choices, the size of the choice sets, and the corresponding number of fixed effects introduce significant computational challenges. In this section, I provide an overview of the assumptions required for estimation and the steps taken to make estimation tractable, deferring further details to Appendix B.1.

I estimate the model separately for each MSA-group, and divide devices into two groups based on whether they are above the median income in the sample MSAs ('higher-income') or below ('lower-income'), where the cutoff is an estimated household income of about \$70,000.<sup>18</sup> Within an MSA-income group, I estimate the model sequentially, which allows me to partition the choice set and only consider alternatives within a given subcategory. Starting with the lowest level of the model (i.e., the choice of an establishment within a subcategory) I estimate the establishment-level fixed effects  $(\hat{\gamma}_j^k)$  and the disutility of driving  $(\hat{\kappa}_t)$ for each subcategory by maximizing the log-likelihood of observing the data on visits to each establishment. The likelihood contribution of any single visit is derived from the choice probabilities  $P_{itj|l_j}$  (Equation 4.4) evaluated at the current parameters.

In estimating the lowest level of the model, I assume that driving times are uncorrelated

 $<sup>^{17}</sup>$ See Appendix B.1 for the formulas for each set of inclusive values.

<sup>&</sup>lt;sup>18</sup>For estimation, I include a buffer region of residents and establishments within 5 miles of the MSA's border. Without this buffer region, neighborhoods close to the border would look mechanically worse than those closer to the center, as many nearby amenities would be outside the estimation sample. I limit preference heterogeneity to just two groups, as the number of group-specific establishment fixed effects quickly becomes infeasible for estimation if I introduce additional types.

with the idiosyncratic component of utility, and, moreover, that the disutility of driving time follows the simple linear form in Equation 4.1. If residents choose to co-locate closer to their favorite establishments (or, alternatively, if establishments locate farther from their most loyal customers towards more marginal customers), the driving times will be endogenous and the estimated coefficients may be biased either up or down (Cao et al., 2024). Similarly, if the true disutility of driving time is non-linear, then the establishment-level fixed effects may be biased in either direction.<sup>19</sup> I show robustness to relaxing each of these assumptions in Section 5.4.

With estimates from the lowest level in hand, I then jointly estimate the upper levels of the model using aggregate data on the number of visits an individual makes to each subcategory. For each individual *i* in quarter *q* at time of week *t*, I observe a vector  $\vec{n}_{itq}$  of total visits to each subcategory. For estimation, I need to take a stance on how often individuals have an opportunity to visit amenities. I assume each individual has two opportunities per day during each time period (i.e. morning, afternoon, and evening).<sup>20</sup> I then maximize the log-likelihood of the data. The log-likelihood contribution of a single  $\vec{n}_{iqt}$  is

$$\ell(\vec{n}_{iqt}) = n_{itq0} \ln\left(P_{it0}\right) + \sum_{l \in \mathcal{L}} n_{itql} \ln(P_{itl})$$

$$(4.6)$$

where  $P_{it0} = 1 - \sum_{l \in \mathcal{L}} P_{itl}$  is the probability of choosing the outside option and  $n_{itq0}$  is the implied number of times the outside option was chosen.

To implement these steps at scale, I use a relatively new computational library from the machine learning world called PyTorch (Paszke et al., 2019). While PyTorch can be used for any optimization problem—including those within reach of more commonly used tools for economics<sup>21</sup>—it is especially well-suited to optimization problems where both the data

<sup>&</sup>lt;sup>19</sup>For example, if the marginal disutility is linear in the log of driving time from (i.e. concave), then the  $\gamma_i$  for establishments farther from most home locations will be biased upwards.

<sup>&</sup>lt;sup>20</sup>I compute the number of choice opportunities an individual has in a quarter using the number of days their device is active that quarter. In a handful of cases, individuals have more observed choices than assumed choice opportunities that quarter, so I use a random sample of their observed choices.

<sup>&</sup>lt;sup>21</sup>For examples of other economics papers using PyTorch, see Lewis, Ozaltun and Zervas (2021) and Du, Kanodia and Athey (2023).

and parameter space are large (e.g., estimation of Large Language Models, like ChatGPT). Its main advantages include automatic differentiation of gradients, seamless integration with GPUs for faster computation, and methods for estimating using larger-than-memory data by iterating over batches of data.

**Parameter estimates.** Appendix Figure B.1 plots the disutility to driving time by time of week. Individuals are more reluctant to drive farther from home than from work, especially during weekday evenings and weekends. Appendix Table B.1 documents the category-time fixed effects ( $\gamma_{mt}$ ) and the degree of independence for each category ( $\rho_m$ ). Subcategories within shops exhibit the most independence from each other, while personal service subcategories exhibit the least. Individuals place relatively more value on restaurants in the afternoons and evenings and more on shops in the morning.

## 4.3 Neighborhood Amenity Quality Index (NAQI)

Using the model framework, I define a 'Neighborhood Amenity Quality Index' (NAQI) that measures the value of each Census block group's access to amenities. For a device of type kliving in Census block group g, the expected value of making a choice at time of week t is given by the usual log-sum formula

$$\mathrm{EV}_{gt}^{k} = \frac{1}{\beta^{k}} \log \left( 1 + \sum_{m \in \mathcal{M}} \exp\left(\gamma_{tm}^{k} + \rho_{m}^{k} \mathrm{IV}_{gtm}\right) \right)$$
(4.7)

where  $IV_{gtm}$  is the inclusive value of each category for someone living in g with no workplace (see Equation B.2) and I divide by  $\beta_k$ , the coefficient on DTEs in Equation 4.5, to convert utility into minutes of driving time. I then compute  $NAQI_g^k$  by summing  $EV_{gt}^k$  over all choice opportunities in a week and normalizing relative to the median neighborhood in each MSA.

A neighborhood will have a higher NAQI if it is closer to establishments that residents value highly. By construction, the NAQI values of a neighborhood are increasing in the number of amenities nearby. However, the extent to which it increases depends on residents' 'love for variety,' which can vary across categories. For example, Appendix Table C.1 documents the variance of each subcategory's idiosyncratic component of utility. Subcategories with higher variance (e.g., full-service restaurants, automobile dealers, and performing arts or spectator sports) offer greater returns for additional establishments than subcategories where the variance is smaller (e.g., general merchandise, grocery stores, dry cleaning, and pharmacies).<sup>22</sup>

# 5 Results

Using the estimated model, I evaluate how preferences for amenities and neighborhoods vary by income, and whether observable characteristics of each are predictive of preferences. I then evaluate the potential welfare consequences of the tailoring of amenities to higher-income residents and the displacement of long-time residents from gentrifying neighborhoods. While I focus on results for average MSA in this section, Appendix Table C.2 documents many of the key results separately by MSA.

## 5.1 Preferences for amenities

For policymakers interested in the consequences of a proposed change to a neighborhood's amenities, a key challenge is identifying preferences for specific amenities using only characteristics that are observable ex-ante. For restaurants, where I can observe detailed characteristics from Yelp, I find that preferences vary by income for observables such as cuisine, price level, and chain size. Figure 3 documents the coefficients from a regression of a restaurant's establishment-specific value  $\gamma_j^k$  (in units of minutes of driving time) on different sets of characteristics, with fixed effects for the MSA. Higher-income residents are willing to drive relatively farther for New American, sushi, and Italian restaurants, and lower-income res-

 $<sup>^{22}</sup>$ Agarwal, Jensen and Monte (2017) find that credit card expenditure shares decline more rapidly with distance in sectors transacted more frequently. In the model presented here, the implied visit shares will gradually decline with distance if the variance of the idiosyncratic component is larger. Consistent with their finding that distance matters least for entertainment establishments and for sellers of durable goods such as cars, I estimate that the variance of the idiosyncratic components is high in these subcategories.

idents place relatively greater value on Mexican and burger restaurants. Similar to Davis et al. (2019), preferences are increasing in the Yelp price level for higher-income residents but generally flat or decreasing for lower-income residents. Lower-income residents also place more value on restaurants that are part of larger chains, while higher-income residents prefer smaller chains (although both place the least value on chains with over 500 establishments).<sup>23</sup>

Many subcategories are also dominated by large brands. Figure 4 plots coefficients from a similar regression of an establishment's value  $(\gamma_j^k)$  on indicators for its affiliation with a subset of chains, with all other chain-affiliated establishments in the subcategory combined into an 'other' group. While higher-income residents prefer Costco to Walmart and lowerincome residents prefer Walmart to Costco, both income groups prefer either option over the many dollar store chains, similar to the findings in Cao (2022) and Cao et al. (2024). Similarly, higher-income residents put relatively more value on Whole Foods than lowerincome residents, but both groups prefer larger chains such as Meijer.<sup>24</sup> This pattern also arises for fitness centers and limited-service restaurants: while some differences between the income groups arise for specific chains, there is a lot of verticality to preferences. Consistent with the intuition for identifying preferences, many chains offering the greatest value have the most visitors and the visitors that travel farther from home (Figure 2).

Even absent more detailed observables, the subcategory of an establishment alone is informative about its potential value to residents. I measure each subcategory's overall contribution to welfare and the correlation of preferences for establishments within a subcategory. To measure a subcategory's welfare contribution, I remove all establishments within the subcategory and compute the average welfare loss for devices in weekly minutes of driv-

 $<sup>^{23}</sup>$ The establishment fixed effects also absorb characteristics of the neighborhood that make them more or less appealing to visit (e.g., easy parking). Appendix Figure C.5 replicates the results in Figure 3 with tract fixed effects. In each case, differences across tracts explain less than half of the differences in preferences for observables.

<sup>&</sup>lt;sup>24</sup>Lower-income residents also prefer Whole Foods over the average grocery store, despite making up a relatively small share of the clientele. Recall that higher valuations can come from either choosing a location frequently or traveling farther to go there. Both higher and lower-income residents travel farther for Whole Foods than the average grocery store chain, with the average lower-income visitor traveling about 40% farther from home (21 minutes for Whole Foods versus 15 minutes for other chains).

ing time. Large subcategories such as malls and full-service restaurants provide substantial value to residents (exceeding 100 minutes/week), while subcategories such as dry cleaning, amusement parks, and bookstores are a much smaller component of welfare (Table 2). Many entertainment subcategories provide more value to higher-income residents, including golf courses, fitness centers, and performing arts or spectator sports. Others provide more value to lower-income residents, including general merchandise or warehouse stores, building material stores, and convenience stores or gas stations. For some categories, the differences between income groups are large: higher-income residents value golf courses and fitness centers about twice as much as lower-income residents.<sup>25</sup>

Higher and lower-income preferences for specific establishments are positively correlated for all subcategories (Table 2). The average correlation is 0.73 when weighted by the total number of stays in each subcategory (0.62 unweighted average). Subcategories such as malls, general merchandise, and parks have especially correlated preferences. Other subcategories especially various personal services and, to a lesser extent, restaurants—have less correlated preferences. In these subcategories, a given establishment is more likely to appeal to only a subset of the income distribution. These cross-subcategory differences may help explain why services and entertainment sectors are more likely to be 'pioneer sectors,' heralding future gentrification in a neighborhood (Behrens et al., 2024).

Preferences for chains are more correlated than preferences for unaffiliated establishments, perhaps because businesses that become chains have broad appeal, or because chains are less flexible in customizing their products and prices to local customers (DellaVigna and Gentzkow, 2019; Klopack, 2024). Excluding chains with at least ten establishments from the four subcategories in Figure 4, the correlation in preferences for establishments drops from 0.76 to 0.34 for general merchandise or warehouse stores, 0.71 to 0.50 for grocery stores, 0.76 to 0.59 for fitness centers, and stays steady at 0.60 for limited-service restaurants.

 $<sup>^{25}</sup>$ Differences in a subcategory's contribution to higher and lower-income welfare are due almost entirely to preferences rather than differences in access from home or work. The welfare contribution of each subcategory changes by at most 4% if a home-work combination for each device is randomly sampled from the entire distribution of home-work combinations.

Independent establishments in these subcategories are more divisive and are more likely to appeal to only a subset of the income distribution.

## 5.2 Preferences for neighborhoods

For each block group, I compute its Neighborhood Amenity Quality Index (NAQI), which is in units of weekly minutes of driving time relative to the median neighborhood in the MSA for each income group.<sup>26</sup> For illustration, Figure 5 maps the NAQI values for the Chicago MSA, zoomed in on Cook County. The most valuable neighborhoods for both higher and lower-income residents are those closer to the downtown area or along the major arterial highways leading into downtown Chicago. While both residents value the amenities of the urban core, differences in preferences arise as we move away from the core towards less dense neighborhoods. Consistent with the income sorting patterns in Chicago, lower-income residents have higher relative values for the amenities of South Chicago, and higher-income

Preferences for neighborhoods are far more correlated than preferences for individual establishments. In the average MSA, the correlation in NAQI is 0.98 (weighted by population), and moving from the 25th to 75th percentile of the NAQI distribution is worth 85 minutes of weekly driving time for a higher-income resident and 87 minutes for a lower-income resident. I decompose differences in the value of neighborhoods' amenity access into saved travel time, increased propensity to visit amenities, and increased quality of amenities visited by simulating choices for devices living in each neighborhood. Saved travel time accounts for 31% of the difference between the 25th and 75th percentile neighborhood for higher-income residents and 33% for lower-income residents. The rest of the difference comes from improved quality of amenities visited, while the propensity to visit amenities is similar at the 25th and 75th percentile neighborhoods for both income groups.

What types of neighborhoods tend to have good access to amenities? The primary

 $<sup>^{26}</sup>$ The NAQI estimates may be valuable inputs for future research and are available to download here.

determinant is proximity to the urban core. The first three columns of Table 3 document the results for univariate regressions of NAQI values on characteristics such as density, median household income, and rent. Each regression includes MSA fixed effects, and the values for non-logged variables correspond to a one standard deviation increase in the covariate. Much as Figure 5 demonstrated for Chicago, the two features most associated with high-quality amenity access are population density and the distance to city hall. A doubling population density is associated with almost a 40-minute/week increase in the NAQI value of a neighborhood for each income group, of which 16-18 minutes come from reduced travel time when visiting amenities. Couture (2016) finds a similar result for restaurants, and shows that almost 40% of the benefit restaurant density comes from reduced travel times. Regressions with either density or distance from city hall have  $R^2$  values of nearly 0.5. For both income groups, the value of a neighborhood's amenity access also increases in rent and education and decreases in household income, fraction white (non-Hispanic), median age, and median income.

Many of the other characteristics are correlated with proximity to the urban core, so the latter three columns control for log density and log distance to city hall. With these controls, neighborhoods with higher income, rent, college graduates, or older residents tend to have higher-quality access to amenities. Despite the levels of positive correlation in NAQI values, some systematic differences emerge. The relationship between a neighborhood's NAQI and its median income, rent, or education is about 25-50% stronger for higher-income residents than lower-income residents, consistent with research on the endogenous response of amenities to local residents (Waldfogel, 2010; Almagro and Dominguez-Iino, 2024).

The returns to living in the urban core are driven by saved driving time and increased variety, but not by better individual establishments. Many of the most highly-valued establishments—such as Costcos, Walmarts, and large malls—are more likely to locate outside the urban core, where there is sufficiently cheap land to have a large footprint. To separate the effects of amenity density from changes in the types of amenities, I recompute the NAQI values after randomizing where each amenity is located, holding fixed the number of amenities in each block group (and setting aside the physical impracticality of swapping a Costco and a restaurant). In this scenario, the unconditional relationship between population density and the NAQI values is 17-19% greater for both income groups.

## 5.3 Tailoring & displacement

Conversations on gentrification often center on two potential channels through which incumbents of the neighborhood may be harmed: amenities may *tailor* to the newcomers and incumbents may be *displaced* by rising rents. In this section, I evaluate the potential for each channel to affect welfare using stylized counterfactuals.

### Tailoring of amenities to higher-income residents.

Amenities that tailor to higher-income residents may amplify welfare inequality by crowding out amenities enjoyed by the lower-income incumbents (Diamond, 2016; Almagro and Dominguez-Iino, 2024). To evaluate the potential magnitude of welfare effects through this channel, I simulate a counterfactual world in which the top 25% of establishments for lowerincome residents within each subcategory are replaced with duplicates of the top 25% of establishments for higher-income residents. The new establishments are set in the same physical location as the ones they replace, but I switch out the establishment-level preference parameters with those corresponding to the new establishments.<sup>27</sup>

I find that this fairly extreme version of tailoring has only a modest effect on welfare. In the average MSA, a higher-income resident's utility increases by 8.2 minutes/week of driving time, and a lower-income resident's utility decreases by 7.1 minutes/week. In a similar exercise, Davis et al. (2019) simulate replacing both restaurants and residents of neighborhoods surrounding majority-Black Harlem with those from majority-white Upper East Side and find that the change in restaurants has only a small effect on welfare relative to

<sup>&</sup>lt;sup>27</sup>The establishment-level fixed effects may, in part, capture preferences for features of the neighborhood in which an establishment is located (e.g., ease of parking). This counterfactual implicitly assumes that such features would be transported along with the establishment.

the social frictions that arise from changing the neighborhood from mostly Black to mostly white. Similar to their findings, the results presented here suggest that the direct effect of amenities tailoring to higher-income residents on spatial inequality is relatively small, though these partial equilibrium effects may be further amplified by general equilibrium channels such as rent increases and residential re-sorting.

#### Displacement of residents from high-rent neighborhoods.

Even absent changes to the types of amenities measured here, incumbents of a gentrifying neighborhood facing rising rents may choose to move to a cheaper neighborhood, with perhaps worse access to amenities and economic opportunities (Newman and Wyly, 2006; Ding, Hwang and Divringi, 2016; Qiang, Timmins and Wang, 2021). I follow Pennington (2021) and Qiang, Timmins and Wang (2021) and identify potentially displaced residents using moves out of neighborhoods with growing rents.<sup>28</sup> First, I identify the set of block groups that experienced at least 25% growth in real rents between 2015-2019 and were in the top within-MSA quartile of average NAQI in 2019. Second, I use data on individual migrations from Infutor to identify all residents who had lived at an address in one of these block groups for at least five years by 2015 and then moved to a cheaper neighborhood sometime between 2015 and 2019.<sup>29</sup> Finally, I compute the average change in NAQI and neighborhood rent for these 'displaced' residents.

The average displaced resident moves to a neighborhood with lower median rents but worse access to amenities.<sup>30</sup> The median block group rent of their new neighborhood is \$469 per month (30%) less than their original neighborhood, and the NAQI value decreases by 40 minutes/week. Relative to even the fairly extreme version of tailoring simulated above,

<sup>&</sup>lt;sup>28</sup>Displacement is often qualitatively defined as involuntary moves caused by forces that "are beyond the household's reasonable ability to control." (Grier and Grier, 1980), such as rising rents, evictions, and natural disasters (Desmond and Shollenberger, 2015). While I cannot separate voluntary from involuntary moves, Qiang, Timmins and Wang (2021) show that lower-income renters are the most likely to leave gentrifying neighborhoods.

<sup>&</sup>lt;sup>29</sup>See Appendix Section A.5 for more details on the Infutor data, which has become a popular source of migration data in the literature (see, e.g., Diamond, McQuade and Qian, 2019; Collinson et al., 2024).

 $<sup>^{30}</sup>$ I average across higher and lower-income NAQI values as I cannot separately identify moves by lower-income versus higher-income residents in the Infutor data.

the risk of displacement is likely the more relevant channel for incumbent residents of a gentrifying neighborhood.

## 5.4 Robustness tests & alternative sources of heterogeneity

In this section, I estimate the sensitivity of key results to a number of assumptions and document heterogeneity by demographics besides income. To limit the computational costs, I estimate all results for a subset of 10 MSAs.<sup>31</sup> The first section of Table 4 document that the results for this sample of MSAs are similar to the overall sample.

Alternative Income measures. The baseline estimates define higher and lower-income using whether a device's estimated income is above or below the median in the sample MSAs. Using whether above/below median income *within-MSA* or discarding the parcel-level estimates in favor of the block group median income from the ACS has little effect on results. Comparing the top and bottom terciles of income reduces the correlation in preferences of POIs by 12% and of neighborhoods by 7% and increases the welfare consequences of tailoring amenities to higher-income residents. As we might expect, preferences in the tails of the income distribution are more polarized than those of residents in the middle tercile.

**Driving time disutility.** I assume in estimation that the driving times to establishments are exogenous to idiosyncratic preferences. In a similar model, Cao et al. (2024) instrument for distances and find that treating distance as exogenous understates the disutility of traveling farther from home, suggesting firms may be endogenously moving farther from their most loyal customers towards more marginal customers. I test how this bias may affect the results here by calibrating the disutility in the model to be 70% higher than the baseline estimate (and re-estimate the rest of the model). I also test the opposite direction of bias, which may occur if residents choose to live near firms for which they have high idiosyncratic

<sup>&</sup>lt;sup>31</sup>The ten MSAs are Chicago-Naperville-Elgin, IL-IN-WI; Miami-Fort Lauderdale-West Palm Beach, FL; Kansas City, MO-KS; Atlanta-Sandy Springs-Roswell, GA; St. Louis, MO-IL; Sacramento–Roseville–Arden-Arcade, CA; Phoenix-Mesa-Scottsdale, AZ; San Antonio-New Braunfels, TX; Seattle-Tacoma-Bellevue, WA; and Minneapolis-St. Paul-Bloomington, MN-WI.

tastes, by calibrating the disutility to be half the baseline disutility. The welfare differences get larger as I decrease the elasticity (and conversely when I increase it), but the magnitudes are small in both cases.

A related concern is whether the disutility of driving time is truly linear. I re-estimate the model using the log of driving time, which reduces the correlation in preferences of POIs by 13%, increases the welfare effects of tailoring amenities to higher-income residents by 6-10 minutes/week, and increases the welfare effects of displacement from gentrifying neighborhoods by 16 minutes/week. Finally, I test sensitivity to my choice of the disutility of driving from home on weekdays as a numeraire by normalizing utility using the average disutility from home across all times of the week and find little change in results.

Subset of trips. The third set of results tests for sensitivity to trip chains and attenuation bias. I first subset to amenity visits that start and end at home (34% of visits) when estimating the lower levels of the model, eliminating trips with multiple stops. This reduces the correlation in both POI and neighborhood preferences and results in a more tightly distributed NAQI distribution. However, much of this effect is likely due to attenuation bias, as estimation on a random one-third subset delivers similar changes relative to the baseline. This leads to a new concern: perhaps the baseline results are also subject to attenuation bias. However, the differences disappear as I increase the sample, and at a two-thirds random sample the results are nearly identical to the full sample.

Alternative sources of heterogeneity. Panel b) of Table 4 documents the results when devices are split by race or education instead of income. The first row splits devices by whether a device lives in a majority non-Hispanic white neighborhood ('white devices,' or WDs) or not (NWDs). The correlation in preferences for amenities and neighborhoods is lower between WD and NWDs than between higher and lower-income devices. Nonetheless, tailoring 25% of establishments to the preference of WDs has only a small effect on utility relative to either the interquartile range of NAQI value or the change in neighborhood NAQI values following potential displacement from a resident's long-time neighborhood. The second row splits devices by whether or not they have a college degree, which I infer from parcel-level characteristics following a similar procedure to income estimation. The results when split by education look similar to when split by income.

# 6 Conclusion

This paper provides new evidence on heterogeneity in preferences for urban amenities across establishments and neighborhoods. Using large-scale GPS data on visits to establishments, I find that, while preferences for specific establishments vary by income, preferences over neighborhoods are highly aligned. Rather than neighborhoods for the rich and others for the poor, cities generally contain neighborhoods sufficiently dense in amenities to offer broad appeal and others with more limited appeal.

These findings have important implications for our understanding of neighborhood change and urban policy. While the entry of specific high-end establishments may signal gentrification (Glaeser, Kim and Luca, 2017; Behrens et al., 2024), it is the broader forces of agglomeration that primarily drive differences in the value of neighborhoods' amenity access. On net, the welfare impacts of tailoring a neighborhood's amenities to higher-income residents are modest compared to the potential effects of residential displacement from gentrifying neighborhoods. Policies that preserve specific local businesses are likely to have only limited effects on the overall value of a neighborhood's amenity access and on spatial sorting by income. Instead, policymakers concerned about the effects of gentrification on the lower-income incumbents of a neighborhood may need to focus more on affordable housing programs and other policies that address displacement risks.

# References

- Agarwal, Sumit, J. Bradford Jensen, and Ferdinando Monte. 2017. "Consumer Mobility and the Local Structure of Consumption Industries." National Bureau of Economic Research w23616.
- Ahlfeldt, Gabriel M., Stephen J. Redding, Daniel M. Sturm, and Nikolaus Wolf. 2015. "The Economics of Density: Evidence From the Berlin Wall." *Econometrica*, 83(6): 2127–2189.
- Allcott, Hunt, Levi Boxell, Jacob Conway, Billy Ferguson, Matthew Gentzkow, and Benjamin Goldman. 2020. "What Explains Temporal and Geographic Variation in the Early US Coronavirus Pandemic?" National Bureau of Economic Research w27965, Cambridge, MA.
- Almagro, Milena, and Tomas Dominguez-Iino. 2024. "Location Sorting and Endogenous Amenities: Evidence from Amsterdam." *Econometrica*, forthcoming.
- Almagro, Milena, Eric Chyn, and Bryan Stuart. 2023. "Urban Renewal and Inequality: Evidence from Chicago's Public Housing Demolitions." National Bureau of Economic Research w30838, Cambridge, MA.
- Athey, Susan, Billy Ferguson, Matthew Gentzkow, and Tobias Schmidt. 2021. "Estimating experienced racial segregation in US cities using large-scale GPS data." *Proceedings of the National Academy of Sciences*, 118(46).
- Athey, Susan, David Blei, Robert Donnelly, Francisco Ruiz, and Tobias Schmidt. 2018. "Estimating Heterogeneous Consumer Preferences for Restaurants and Travel Time Using Mobile Location Data." AEA Papers and Proceedings, 108: 64–67.
- Barwick, Panle Jia, Shanjun Li, Andrew R. Waxman, Jing Wu, and Tianli Xia. 2024. "Efficiency and Equity Impacts of Urban Transportation Policies with Equilibrium Sorting." *American Economic Review*, forthcoming.
- Baum-Snow, Nathaniel, and Daniel Hartley. 2020. "Accounting for central neighborhood change, 1980–2010." Journal of Urban Economics, 117: 103228.
- Bayer, Patrick, and Christopher Timmins. 2007. "Estimating Equilibrium Models of Sorting Across Locations." *The Economic Journal*, 117(518): 353–374.
- Bayer, Patrick, and Robert McMillan. 2012. "Tiebout sorting and neighborhood stratification." Journal of Public Economics, 96(11): 1129–1143.
- Behrens, Kristian, Brahim Boualam, Julien Martin, and Florian Mayneris. 2024. "Gentrification and Pioneer Businesses." *The Review of Economics and Statistics*, 106(1): 119–132.
- Berkes, Enrico, and Ruben Gaetani. 2023. "Income Segregation and the Rise of the Knowledge Economy." American Economic Journal: Applied Economics, 15(2): 69–102.
- Birant, Derya, and Alp Kut. 2007. "ST-DBSCAN: An algorithm for clustering spatial-temporal data." Data & Knowledge Engineering, 60(1): 208-221.
- Brummet, Quentin, and Davin Reed. 2021. "The Effects of Gentrification on the Well-Being and Opportunity of Original Resident Adults and Children." Working paper 3421581.
- Caetano, Gregorio, and Vikram Maheshri. 2019. "Gender segregation within neighborhoods." *Regional Science and Urban Economics*, 77: 253–263.
- Cao, Yue. 2022. "The Welfare Impact of Dollar Stores." Working paper.

- Cao, Yue, Judith A Chevalier, Jessie Handbury, Hayden Parsley, and Kevin R Williams. 2024. "Distributional Impacts of the Changing Retail Landscape." Working paper.
- Chang, Serina, Emma Pierson, Pang Wei Koh, Jaline Gerardin, Beth Redbird, David Grusky, and Jure Leskovec. 2021. "Mobility network models of COVID-19 explain inequities and inform reopening." *Nature*, 589(7840): 82–87. Publisher: Nature Publishing Group.
- Chen, M. Keith, Kareem Haggag, Devin G. Pope, and Ryne Rohla. 2022. "Racial Disparities in Voting Wait Times: Evidence from Smartphone Data." *The Review of Economics and Statistics*, 104(6): 1341–1350.
- Chen, Tianqi, and Carlos Guestrin. 2016. "XGBoost: A Scalable Tree Boosting System." In Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. KDD '16, 785–794.
- Collinson, Robert, John Eric Humphries, Nicholas Mader, Davin Reed, Daniel Tannenbaum, and Winnie van Dijk. 2024. "Eviction and Poverty in American Cities." *The Quarterly Journal of Economics*, 139(1): 57–120.
- Cook, Cody, and Pearl Z Li. 2023. "Value Pricing or Lexus Lanes? The Distributional Effects of Dynamic Tolling." Working paper.
- Cook, Cody, Lindsey Currier, and Edward Glaeser. 2024. "Urban mobility and the experienced isolation of students." *Nature Cities*, 1(1): 73–82.
- **Couture, Victor.** 2016. "Valuing the Consumption Benefits of Urban Density." Working paper.
- Couture, Victor, and Jessie Handbury. 2020. "Urban Revival in America, 2000 to 2010." Journal of Urban Economics, 119.
- **Couture, Victor, and Jessie Handbury.** 2023. "Neighborhood Change, Gentrification, and the Urbanization of College Graduates." *Journal of Economic Perspectives*, 37(2): 29–52.
- Couture, Victor, Cecile Gaubert, Jessie Handbury, and Erik Hurst. 2023. "Income Growth and the Distributional Effects of Urban Spatial Sorting." *The Review of Economic Studies*, 91(2): 858–898.
- Couture, Victor, Jonathan I. Dingel, Allison Green, Jessie Handbury, and Kevin R. Williams. 2022. "JUE Insight: Measuring movement and social contact with smartphone data: a real-time application to COVID-19." *Journal of Urban Economics*, 127: 103328.
- **Davis, Donald R., Jonathan I. Dingel, Joan Monras, and Eduardo Morales.** 2019. "How Segregated Is Urban Consumption?" *Journal of Political Economy*, 127(4).
- **DellaVigna, Stefano, and Matthew Gentzkow.** 2019. "Uniform Pricing in U.S. Retail Chains." *The Quarterly Journal of Economics*, 134(4): 2011–2084.
- **Desmond, Matthew, and Tracey Shollenberger.** 2015. "Forced Displacement From Rental Housing: Prevalence and Neighborhood Consequences." *Demography*, 52(5): 1751–1772.
- Diamond, Rebecca. 2016. "The Determinants and Welfare Implications of US Workers' Diverging Location Choices by Skill: 1980–2000." American Economic Review, 106(3): 479–524.

- Diamond, Rebecca, Tim McQuade, and Franklin Qian. 2019. "The Effects of Rent Control Expansion on Tenants, Landlords, and Inequality: Evidence from San Francisco." *American Economic Review*, 109(9): 3365–3394.
- Ding, Lei, Jackelyn Hwang, and Eileen Divringi. 2016. "Gentrification and residential mobility in Philadelphia." *Regional Science and Urban Economics*, 61: 38–51.
- Dolfen, Paul, Liran Einav, Peter J. Klenow, Benjamin Klopack, Jonathan D. Levin, Larry Levin, and Wayne Best. 2023. "Assessing the Gains from E-Commerce." *American Economic Journal: Macroeconomics*, 15(1): 342–370.
- Du, Tianyu, Ayush Kanodia, and Susan Athey. 2023. "Torch-Choice: A PyTorch Package for Large-Scale Choice Modelling with Python." Working paper.
- Ellen, Ingrid Gould, and Katherine O'Regan. 2010. "Crime and urban flight revisited: The effect of the 1990s drop in crime on cities." *Journal of Urban Economics*, 68(3): 247–259.
- Ellen, Ingrid Gould, and Lei Ding. 2016. "Advancing Our Understanding of Gentrification." *Cityscape*, 18(3).
- Foster, Kevin, Claire Greene, and Joanna Stavins. 2020. "The 2019 Survey of Consumer Payment Choice: Summary Results." Federal Reserve Bank of Atlanta.
- **Glaeser, Edward, Hyunjin Kim, and Michael Luca.** 2017. "Nowcasting the Local Economy: Using Yelp Data to Measure Economic Activity." National Bureau of Economic Research w24010, Cambridge, MA.
- Glaeser, Edward L., Jed Kolko, and Albert Saiz. 2001. "Consumer city." Journal of Economic Geography, 1(1): 27–50.
- Glass, Ruth. 1964. "Aspects of change." In *The gentrification debates*. 19–29. Routledge.
- Gordinier, Jeff. 2016. "South Bronx Gets High-End Coffee; Is Gentrification Next?" New York Times.
- Grier, George, and Eunice Grier. 1980. "Urban displacement: a reconnaissance." In *Back to the City.* 252–268. Pergamon.
- Guerrieri, Veronica, Daniel Hartley, and Erik Hurst. 2013. "Endogenous gentrification and housing price dynamics." *Journal of Public Economics*, 100: 45–60.
- Gupta, Arpit, Stijn Van Nieuwerburgh, and Constantine Kontokosta. 2022. "Take the Q train: Value capture of public infrastructure projects." *Journal of Urban Economics*, 129: 103422.
- Handbury, Jessie. 2021. "Are Poor Cities Cheap for Everyone? Non-Homotheticity and the Cost of Living Across U.S. Cities." *Econometrica*, 89(6): 2679–2715.
- Handbury, Jessie, and David E. Weinstein. 2015. "Goods Prices and Availability in Cities." *The Review of Economic Studies*, 82(1 (290)): 258–296.
- Klopack, Ben. 2024. "One size fits all? The value of standardized retail chains." The RAND Journal of Economics, 55(1): 55–86.
- Kolomatsky, Michael. 2020. "Whole Foods and Trader Joe's Downstairs, Higher Rent Upstairs." New York Times.
- Lester, T. William, and Daniel Hartley. 2014. "The long term employment impacts of gentrification in the 1990s." *Regional Science and Urban Economics*, 45(C): 80–89. Publisher: Elsevier.
- Lewis, Greg, Bora Ozaltun, and Georgios Zervas. 2021. "Maximum Likelihood Estimation of Di erentiated Products Demand Systems." Working paper.

- McFadden, Daniel. 1973. "Conditional Logit Analysis of Qualitative Choice Behavior." *Frontiers in Econometrics*, 105–142.
- McKinnish, Terra, Randall Walsh, and T. Kirk White. 2010. "Who gentrifies lowincome neighborhoods?" Journal of Urban Economics, 67(2): 180–193.
- Meltzer, Rachel, and Pooya Ghorbani. 2017. "Does gentrification increase employment opportunities in low-income neighborhoods?" *Regional Science and Urban Economics*, 66(C): 52–73.
- Miyauchi, Yuhei, Kentaro Nakajima, and Stephen J. Redding. 2022. "Consumption Access and the Spatial Concentration of Economic Activity: Evidence from Smartphone Data." Working paper.
- Newman, Kathe, and Elvin K. Wyly. 2006. "The Right to Stay Put, Revisited: Gentrification and Resistance to Displacement in New York City." Urban Studies, 43(1): 23–57.
- Oh, Ryungha, and Jaeeun Seo. 2023. "What Causes Agglomeration of Services?" Paszke, Adam, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gre-
- gory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. "PyTorch: An Imperative Style, High-Performance Deep Learning Library." In Advances in Neural Information Processing Systems. Vol. 32. Curran Associates, Inc.
- **Pennington, Kate.** 2021. "Does Building New Housing Cause Displacement?: The Supply and Demand Effects of Construction in San Francisco." Working paper SSRN Scholarly Paper 3867764, Rochester, NY.
- Phillips, David C. 2020. "Measuring Housing Stability With Consumer Reference Data." Demography, 57(4): 1323–1344.
- Qiang, Ashley J, Christopher Timmins, and Wen Wang. 2021. "Displacement and the Consequences of Gentrification." Working paper.
- Raleigh, Helen. 2017. "Gentrification Provokes a Coffee Clash in Denver's Five Points." *Wall Street Journal.*
- Reardon, Sean F., and Kendra Bischoff. 2011. "Income Inequality and Income Segregation." *American Journal of Sociology*, 116(4): 1092–1153. Publisher: The University of Chicago Press.
- Redding, Stephen J, and Esteban Rossi-Hansberg. 2017. "Quantitative Spatial Economics." Annual Review of Economics, 9:21–58.
- **Relihan, Lindsay E.** 2024. "Clicks and Bricks: The Complementarity of Online Retail and Urban Services." Working paper.
- Roback, Jennifer. 1982. "Wages, Rents, and the Quality of Life." Journal of Political Economy, 90(6): 1257–1278.
- **Rosen, Sherwin.** 1981. "The Economics of Superstars." *The American Economic Review*, 71(5): 845–858.
- Small, Kenneth A., and Harvey S. Rosen. 1981. "Applied Welfare Economics with Discrete Choice Models." *Econometrica*, 49(1): 105–130.
- Smith, Rosa Inocencio. 2016. "When a Grocery Store Means Gentrification." The Atlantic.

- Su, Yichen. 2022. "The Rising Value of Time and the Origin of Urban Gentrification." American Economic Journal: Economic Policy, 14(1): 402–439.
- **Tiebout, Charles M.** 1956. "A Pure Theory of Local Expenditures." Journal of Political Economy, 64(5): 416–424.
- **Train, Kenneth E.** 2009. *Discrete Choice Methods with Simulation*. Cambridge University Press.

Ukueberuwa, Mene. 2020. "Gentrification Is America's Best Hope." Wall Street Journal.

- Vigdor, Jacob L. 2002. "Does Gentrification Harm the Poor?" Brookings-Wharton Papers on Urban Affairs, 2002(1): 133–182. Publisher: Brookings Institution Press.
- Waldfogel, Joel. 2003. "Preference Externalities: An Empirical Study of Who Benefits Whom in Differentiated-Product Markets." *The RAND Journal of Economics*, 34(3): 557–568.
- Waldfogel, Joel. 2010. "Who Benefits Whom in the Neighborhood? Demographics and Retail Product Geography." In Agglomeration economics. University of Chicago Press.

# Tables

Table 1: GPS sample

		Most populous MSAs				
	All	Los Angeles	Chicago	Dallas	Houston	
Num. devices	7,748,661	640,584	483,809	$562,\!621$	504,025	
Avg. total stays per quarter	94.88	91.59	96.73	100.31	99.29	
	[86.0]	[85.92]	[86.52]	[89.93]	[89.21]	
Avg. amenity visits per quarter	11.55	10.86	10.95	12.65	12.06	
	[18.56]	[17.68]	[17.02]	[19.22]	[19.09]	
Avg. home stays per quarter	47.91	47.34	49.08	48.96	48.45	
	[38.73]	[39.27]	[39.15]	[39.05]	[38.76]	
Avg. work stays per quarter	11.95	11.33	12.21	12.97	12.71	
	[17.58]	[17.60]	[17.46]	[18.36]	[18.4]	

*Notes:* This table documents the sample size and the number of stays observed for the average devicequarter, for the full sample and for the four most populous MSAs. Standard deviations are presented in brackets. All averages are weighted by the device's sample weight.

	Contribution to welfare (per-person minutes/week)		$\begin{array}{c} \text{Correlation} \\ (\gamma_j^H,  \gamma_j^L) \end{array}$
	Н	L	
Restaurants			
Full-service	96.18	112.40	0.73
Limited-service	29.83	29.24	0.60
Cafes; snacks	22.23	16.01	0.59
Drinking places	16.73	15.24	0.56
Shops			
Malls	226.61	180.61	0.87
General merch., warehouse clubs	36.33	56.72	0.76
Groceries, beer/wine/liquor stores	26.14	32.01	0.71
Clothing, shoes, jewelry, leather goods	27.24	29.77	0.70
Gas stations, convenience	14.25	26.94	0.64
Department	16.57	18.25	0.67
Building materials, gardening	14.72	19.49	0.70
Furniture, appliances, electronics	10.58	11.11	0.62
Automobile dealers	9.72	11.49	0.66
Sporting goods, hobby, music	8.59	8.95	0.63
Pharmacies	7.57	9.75	0.60
Beauty, glasses, personal care	6.94	8.09	0.66
Books, office supplies	4.30	4.74	0.62
Personal services			
Hospitals, health clinics	43.51	42.66	0.73
Salons, barbers	21.73	16.93	0.45
Religious organizations	16.71	13.28	0.40
Banks	10.35	9.10	0.56
Dentists	6.90	5.01	0.40
Car maintenance	2.33	2.33	0.28
Drycleaning, laundry	0.78	0.81	0.42
Entertainment			
Parks	86.70	70.47	0.75
Fitness centers	42.05	24.68	0.66
Golf courses	31.21	13.68	0.59
Performing arts and spectator sports	20.13	12.49	0.70
Movie theaters	12.28	9.07	0.74
Gambling	2.87	14.67	0.60
Other amusement, recreation	7.02	5.01	0.67
Libraries, museums, zoos, gardens	5.45	3.92	0.61
Amusement parks	4.00	3.29	0.76

# Table 2: Subcategory correlations & welfare

*Notes:* This table documents the correlation between establishment-level preferences within each subcategory and each subcategory's contribution to welfare. Welfare contribution is computed as the change in expected utility when all establishments in the subcategory are removed. Subcategories within a category are ordered by their total welfare contribution.

	Univ	ariate	Controls f & city h	for density nall dist.
	Н	L	Н	L
Log population density	37.91	39.05		
	(SE: 0.21)	(SE: 0.21)		
	[R2: 0.47]	[R2: 0.48]		
Log miles to city hall	-70.05	-73.13		
	(SE: 0.31)	(SE: 0.31)		
	[R2: 0.48]	[R2: 0.51]		
Median income	-1.58	-4.63	10.78	7.81
	(SE: 0.25)	(SE: 0.25)	(SE: 0.18)	(SE: 0.18)
	[R2: 0.01]	[R2: 0.01]	[R2: 0.62]	[R2: 0.64]
Median rent	7.56	5.38	15.32	12.28
	(SE: 0.3)	(SE: 0.3)	(SE: 0.28)	(SE: 0.28)
	[R2: 0.02]	[R2: 0.02]	[R2: 0.61]	[R2: 0.63]
Frac. college grad	15.67	11.84	12.89	8.82
	(SE: 0.25)	(SE: 0.25)	(SE: 0.16)	(SE: 0.16)
	[R2: 0.05]	[R2: 0.03]	[R2: 0.63]	[R2: 0.64]
Frac. white-alone	-20.54	-23.28	1.42	-0.85
	(SE: 0.25)	(SE: 0.25)	(SE: 0.18)	(SE: 0.17)
	[R2: 0.09]	[R2: 0.1]	[R2: 0.6]	[R2: 0.63]
Median age	-13.81	-15.33	4.66	3.52
	(SE: 0.28)	(SE: 0.28)	(SE: 0.21)	(SE: 0.21)
	[R2: 0.04]	[R2: 0.05]	[R2: 0.6]	[R2: 0.63]

Table 3∙	Relationship	hetween	NAOI	and	neighborhood	characteristics
Table 9.	riciationship	DCUWCCII	TATAT	and	neignoornoou	. Characteristics

*Notes:* This table documents results from a series of regressions of the block group level NAQI on neighborhood characteristics. Population and demographics data are from the 2019 5-year ACS. The distance to city hall is crow-flies, based on the city hall location reported on Google Maps for the largest city in each MSA. Regressions include fixed effects for the MSA. For non-log covariates, the coefficients are standardized to correspond to a one standard deviation increase.

			NA	AQI		Counterfactuals	
	Corre	elation	p75	-p25	Tailoring		Displacement
Panel a: robustness	POIs	NAQI.	Н	L	Н	L	All
Baseline							
All MSAs	0.73	0.98	82.5	83.2	8.19	-7.11	-39.7
MSAs for robustness	0.74	0.98	85.7	82.1	8.80	-8.04	-38.7
Income measures							
Block group median income (within-MSA)	0.74	0.98	86.9	83.8	8.56	-6.58	-38.9
Above/below median (within-MSA)	0.74	0.98	85.8	83.3	9.28	-7.09	-38.7
Top/bottom tercile (within-MSA)	0.63	0.92	84.8	75.2	13.32	-9.69	-37.7
Subset of trips							
Home-amenity-home trips	0.66	0.94	74.4	66.6	7.80	-6.26	-34.8
Random $1/3$ subset	0.69	0.85	72.6	67.3	7.70	-7.62	-33.2
Random $2/3$ subset	0.71	0.98	86.0	81.3	8.83	-8.72	-38.7
Driving times							
Increase disutility (Cao et al., 2024)	0.80	0.94	71.4	72.4	8.79	-6.15	-33.7
Decrease disutility	0.75	0.98	98.5	96.6	10.46	-13.46	-40.8
Log driving time	0.64	0.97	84.9	80.3	14.4	-17.8	-54.3
Normalize by average disutility	0.74	0.98	96.7	91.5	9.54	-9.38	-41.7
	Corre	elation	p75	-p25	Tailor		Displace
Panel b: heterogeneity	POIs	NAQI	(1)	(2)	(1)	(2)	All
WD (1) v. NWD (2)	0.69	0.88	85.4	67.1	8.02	-4.30	-36.6
College $(1)$ v. no college $(2)$	0.72	0.98	82.3	84.1	8.45	-7.18	-38.2

Table 4: Robustness tests & additional sources of heterogeneity

*Notes:* This table documents the results for versions of the model estimated either for robustness (panel a) or to document heterogeneity along different dimensions (panel b). The first two columns document the correlation in establishment-level preferences subcategories (weighted by the total number of visits in each subcategory) and NAQI values. The second two document the interquartile range for the NAQI values, which are in units of minutes of driving time per week. The tailoring counterfactual document the effects (in minutes/week of value) of tailoring 25% of establishments to the preferences of higher-income residents. The final column documents the average change in NAQI value for long-time residents of gentrifying neighborhoods that moved to a cheaper neighborhood between 2015-2019. Panel b) documents results for white devices (WDs) and non-white devices (NWD), which are defined by whether a device's home Census block was at least 50% white non-Hispanic in the 2010 Census. College and non-college labels are estimated based on a device's home parcel; see Appendix Section A.2 for more details.

# **Figures**



Figure 1: Neighborhood amenities: consumption and access

(a) Access to amenities

*Notes:* The first panel plots the number of establishments of different categories within a 10-minute drive of a device's home or work block group and the second plots the number of visits to amenities of each category. Each point is the coefficient from a regression of the inverse hyperbolic sine of the outcome on indicators for income quartile using device-quarter level of data (I use the inverse hyperbolic sine to handle zeros, although I refer to it as 'log' for brevity). Income quartile cutoffs are based on the distribution of incomes in the 30 MSAs in the sample. Gray bars represent 95% confidence intervals, with standard errors clustered by home or work block group for panel (a) and by device for panel (b).



Figure 2: Heterogeneity by chain: # visits, driving time, and income

*Notes:* For establishments associated with specific chains, this figure documents the number of visits per establishment and the average income and minutes from home of their visitors. Minutes from home and the number of visitors are standardized by taking the z-score across chains within the same subcategory. The gray dots are other chains in the same subcategory with at least 25 establishments.



Figure 3: Relative value of restaurant characteristics

(a) Cuisines

Notes: This figure documents the relationship between the restaurants' estimate value for each income group and different observables. Each point corresponds to a coefficient of the establishment-level fixed effects  $(\gamma_j^k)$  on characteristics of the establishment with controls for the MSA. Gray bars represent 95% confidence intervals.





(a) General merchandise

Notes: This figure documents the relationship between an establishment's estimated value for each income group and its chain. Each point corresponds to a coefficient of the establishment-level fixed effects  $(\gamma_j^k)$  on characteristics, with fixed effects for the MSA. The data are subset to just chain establishments, and the holdout group is 'other' chains. Gray bars represent 95% confidence intervals.



Figure 5: Neighborhood Amenity Quality Index: Chicago

(a) Higher-income

(b) Lower-income

*Notes:* This figure illustrates the estimated block group level NAQI values for the Chicago MSA, zoomed in to Cook County (outlined in black). NAQI values for each group correspond to minutes of weekly driving time relative to the median neighborhood for that income group.

# Appendices

A	Data appendix	<b>2</b>
	A.1 GPS location data	2
	A.2 Estimating individual income	3
	A.3 Evaluating sample quality	8
	A.4 Trip chains	13
	A.5 Identifying 'displaced' residents in Infutor	13
В	Model appendix	14
	B.1 Estimation details	14
	B.2 Parameter estimates	15
С	Supplementary exhibits	18

# A Data appendix

## A.1 GPS location data

GPS location data allows researchers to follow GPS-enabled devices (primarily smartphones) as they move around their city. The data used in this paper is provided by Replica, a data platform company that purchases raw data from multiple aggregators (e.g., SafeGraph), which they use to build tools for city planners, transportation agencies, and other policy-makers. In this section I describe the processes for cleaning raw GPS location data, assigning homes and works to each device, and identifying visits to amenities.

### A.1.1 Pre-processing: from pings to stays

The raw GPS data consists of 'pings,' which can be triggered either by a user opening specific apps or via apps that share location data in the background.<sup>32</sup> Each ping includes the coordinates, timestamp, and a unique device identifier. The ping-level data are large; the 2019 data include over 700 billion pings.

To reduce the size of the data while retaining the core information, Replica combines individual pings that are close together in time and space into a 'stay' using a variant of the ST-DBSCAN clustering algorithm (Birant and Kut, 2007). Pings while traveling are discarded, and the minimum stay length is five minutes. The coordinates of a stay are the centroid of all pings assigned to the same cluster by ST-DBSCAN, and pings that occur in transit between stays are discarded. Each stay includes the coordinates, enter time, exit time, and unique device identifier.

#### A.1.2 Assigning home and work locations

For each day a device appears in the data, I use stays from the 28 days leading up to a given date to assign a most likely home and—where applicable—work location as of that date.

Home assignment consists of three steps. First, I count all overnight (12-5am) stays by a device in an H3 bound of resolution 9. H3 is a hierarchical spatial index, similar to geohashes. Bounds of resolution 9 cover approximately 100 square meters. Second, I label the most common H3 bound as the home bound, so long as it has at least five overnight dates. Finally, I identify the exact coordinates of the home by taking the centroid of all overnight stays within the home bound.

Assigning workplaces is more complicated, as work shifts can occur at all hours of the day and, for many occupations, may not involve a single location (e.g., plumbers, taxi drivers, and appraisers). I label the non-home H3 bound with the most daytime (6am-8pm) stays of over 2 hours as the work bound of a device and define the work coordinates as the centroid of all stays in this bound.

To determine the home and work locations at the device-quarter level, I use the most commonly assigned home and work across the quarter. I exclude any device-quarters for

<sup>&</sup>lt;sup>32</sup>Prior to 2021, collecting location data in the background was common. Following changes to iOS and Android that prompted users to repeatedly opt-in to background location tracking and limited access to unique device identifiers, GPS data quality declined sharply, and many vendors stopped selling the data.

which there is no reliable home location. Of device-quarters with a home assigned, 67% are assigned a work location.<sup>33</sup>

#### A.1.3 Identifying visits to amenities

I match stays to Points of Interest (POIs) by defining a stay as occurring at a POI if its coordinates fall within the polygon describing its footprint during its operating hours. These polygons are not always disjoint, so a single stay may match multiple POIs. There are three primary cases when polygons may overlap. First, when a POI is entirely enclosed in another and does not have its own polygon, such as a store in a large mall. In this case, I assign the 'parent' POI (i.e. the mall). Second, POIs within a 'parent' polygon may have their own individual polygons in the data (e.g., a shopping center where stores are more easily separated). In this case, the coordinates will match both the parent and the individual POI's polygon, but the data are sufficient to discern which POI within the parent POI was visited, so I assign the individual POI. Finally, polygons may overlap due to being stacked in multi-story buildings. In this case, I randomly assign a polygon among the matching set. Of the stays matched to at least one Safegraph POI, 76.3% either match to a single POI or a single non-parent POI, 9.8% are matched to a parent POI, and 13.9% are randomly assigned a POI from the set of matches. I exclude visits to POIs at a device's home or work location.

I classify each POI into mutually exclusive subcategories of amenities using SafeGraph's assigned NAICS code. Table A.1 documents the NAICS code(s) used for each subcategory. One exception is malls, which do not have a corresponding NAICS code but are the 'parent' POI for many smaller establishments. I use SafeGraph's 'Lessors of Real Estate' category to identify Malls. Finally, I drop establishments with fewer than five observed visits. Table A.1 documents the subcategories, NAICS codes, and the number of POIs.

## A.2 Estimating individual income

I use parcel-level data to predict the household income of devices residing at that parcel and whether they are above the median income for the sample MSAs ('higher-income'). At a high level, the procedure is as follows:

- 1. Match each device to a home parcel using data on parcel polygons and CoreLogic parcel characteristics
- 2. Use data on historical transactions to estimate the 2019 market value of all houses
- 3. Estimate the relationship between household income and house characteristics in the 2019 5-year ACS, then predict household income for each CoreLogic parcel
- 4. Bayesian-update each estimate using the distribution of home block group income in the 2019 5-year ACS tabulations

Each step is described in more detail below.

<sup>&</sup>lt;sup>33</sup>This method for assigning works has drawbacks and cannot identify workplaces for occupations that do not involve a static location (e.g., Uber drivers and mail carriers) or for workers that work night shifts.

Sub-category name	NAICS codes	# POIs	# visits
Restaurants			
Full-service	722511	183985	6627892
Limited-service	722513	87409	2009523
Cafes; snacks	722515	57489	1326584
Drinking places	722410	27808	1235171
Shops			
Malls	N/A	15636	39848901
General merch., warehouse clubs	4523	13987	2726402
Gas stations, convenience	4471,  445120	55800	2300789
Groceries, beer/wine/liquor stores	445110, 4452, 4453	58997	2110776
Clothing, shoes, jewelry, leather goods	4481, 4482, 4483	73791	1587154
Building materials, gardening	4441, 4442	23345	1231001
Automobile dealers	4411, 4412	29561	1191595
Furniture, appliances, electronics	4421, 4422, 4431	48354	1015824
Beauty, glasses, personal care	446120, 446130, 446190	24051	884172
Pharmacies	446110	22730	858123
Sporting goods, hobby, music	4511	31993	621240
Department	4522	6123	570976
Books, office supplies	4512, 4532	17777	314200
Personal Services			
Hospitals, health clinics	6221, 6214	18020	5036359
Salons, barbers	8121	180353	3055200
Religious organizations	8131	69651	1906581
Banks	5221, 5223, 5231	69776	1297590
Dentists	6212	67942	896281
Car maintenance	8111	88137	549098
Drycleaning, laundry	8123	7691	133491
Entertainment			
Parks	712190	47122	8955796
Golf courses	713910	5426	3143206
Fitness centers	713940	48776	2300442
Performing arts and spectator sports	7111, 7113, 7112	4018	1280980
Other amusement, recreation	713990, 713950	11118	580333
Gambling	7132	749	558625
Amusement parks	7131	2725	521346
Movie theaters	512131	4001	502212
Libraries, museums, zoos, gardens	712110, 712130, 519120	8305	363843

Table A.1: Amenity subcategories

*Notes:* This documents the categories and subcategories of amenities, the NAICS codes used to identify each subcategory, and the number of establishments and (unweighted) visits in the sample. NAICS codes are provided by SafeGraph. To identify malls, I use SafeGraph's category 'Lessors of Real Estate.'

#### A.2.1 Matching devices to houses

I match each device to residential parcel polygons provided by LandGrid by spatially joining each set of home coordinates to parcel data. I use a strict match and only include those home coordinates that fall within the bounds of a residential parcel.

Next, I match LandGrid parcels to housing characteristics provided by CoreLogic. Core-Logic has two main data sources: characteristics of a given house from tax assessments filed with local governments ("Tax") and details of each sale of a house from the deed filed during the sale ("Deed"). Most of the Tax data is from the 2018 tax year, while the Deeds extend through June 2019. To match Landgrid parcels to CoreLogic details, I use the coordinates provided by CoreLogic and again spatially join to the polygon from Landgrid.

In total, 73.2% of device homes match to a parcel with CoreLogic characteristics. The majority of the remaining devices were not successfully matched to any residential parcel. This can happen, for example, if their 'home' coordinates were on the street because the home location identification was noisy.

#### A.2.2 Estimating market value of homes

To predict the value of a given device's home, I build a model for estimating the 2019 market value of all housing units in the US (not just those with a matched GPS device). The model is trained on data on actual housing sales from CoreLogic, then used to predict out-of-sample using the location and characteristics of all homes, even those for which I never observe a sale. I focus on predicting the value of single-family houses, townhouses, and condos; large rental apartment buildings are, for now, set aside. The sample includes 104.1 million properties.

To build a training sample, I match each house in the Tax data to any sales between 2010-2019 in the Deed data. I restrict the sales data to arms-length transactions and remove all homes purchased by an owner with 'LLC' in the name, as they often involve purchases of many units that have sale prices corresponding to the total purchase rather than each unit. Finally, I further restrict to houses sold at prices between \$10,000 and \$25,000,000. The final data include 28 million sales.

The Tax data include many characteristics of the property, including: number of bedrooms and bathrooms; living, garage, basement, and land square footage; the style of the building, when it was built, and whether refurbished; a measure of quality (e.g. 'Fair'); what type of view the property has (e.g., 'Mountains') and other 'location influences'; whether the property has a pool; the type of air conditioning, heating, fuel, framing, walls, sewage, water, roof, and floor; and the precise location. Some of these variables are frequently missing – in all such cases, I replace the missing values with an indicator for missing data. I also discretize continuous variables, such as living square footage, into deciles.

To predict the market value of all homes, I train a deep neural net (DNN) on 80% of all sales, setting aside the remaining 20% to evaluate model performance. First, I turn all of the characteristics of the property and the time of the sale into features. Many of these characteristics are high dimensional; for example, there are 87 unique styles of houses and 287 types of exterior walls. For sale year, sale month, property type, decade built, decile of living square feet, and whether there is a pool, I encode the values using dummy variables. For the remaining characteristics described above and each property's Census tract, I use embedding vectors, which are an alternative to dummy variables. Rather than mapping a unique value to a vector of 0s and a single 1, they map each unique value to a vector of continuous values whose weights are learned through the model training.

For model architecture, I use five hidden linear layers with rectified linear units (ReLUs) as activation functions between each layer and batch normalization to increase stability. I apply a sigmoid function to the output to restrict estimates to be between 70% of the lowest observed sale price and 130% of the highest. The model is built using PyTorch and run on GPUs. I train the model separately for each state, using an Adam optimizer, a maximum of 100 epochs (i.e. complete runs through the data), and batch sizes of 256 sales.

Figure A.1 provides a visualization of how well the model performs relative to a linear regression in California.<sup>34</sup> A model that could perfectly separate houses into deciles would be a diagonal line of dark squares. In general, the DNN is far closer to this ideal than the linear regression and is only infrequently off by more than a decile. On the full sample, the DNN has a root mean squared error of 0.64 in the holdout sample, and correctly assigns the quartile of sales prices for 79% of sales (compared to 25% if quartiles were randomly assigned. The results are similar within-sample, suggesting the model is not overfit.

Finally, for all 104.1 million properties in the Tax data—even those that did sell—I predict the market value based on the property characteristics and assuming a June 2019 sale date.



Figure A.1: Deciles of true v. estimated sale amounts: California

*Notes:* These figure shows the out-of-sample quality of different models for California by highlighting the percent of predicted sale value deciles that match the actual sale value.

#### A.2.3 Estimating whether above median income

I use data from the 2019 5-year ACS Public Use Microdata Sample (PUMS) to model the relationship between housing characteristics and income. I then use this model to predict whether a device is above the median income in the sample MSAs (\$69,888) using the

<sup>&</sup>lt;sup>34</sup>For the linear regression, features are encoded as dummy variables instead of embeddings, and for computational purposes, zip codes are used rather than the more granular Census tracts.

characteristics of the parcel in which they reside as well as the overall income distribution of their block group.

I first harmonize the ACS and CoreLogic characteristics, such that a model can be trained on the ACS data and evaluated on CoreLogic data. I use up to four characteristics of each home: its location (Public Use Microdata Area), the number of units in the building, the decade built, and its estimated home value. In the ACS, households self-report their home value ('valueh') or the rent they pay ('rentgrs'). In CoreLogic, I observe only estimated home value, not whether the landlord or renters occupy the parcel. To address this, I use the within-MSA decile of home value or gross rent (whichever is available) instead of the actual estimated market value. This approach is equivalent to simply using the decile of home value if rents are set as a MSA-wide multiplier on the home value.

I use a common machine learning classifier, XGBoost, to estimate whether a household is above median income in the ACS data (Chen and Guestrin, 2016). I train three versions of this model, for varying levels of data availability: 1) using all four characteristics described above, 2) using just home value/rent deciles and the Public Use Microdata Area (PUMA), and 3) using just the decade built, units, and PUMA. Each feature enters as a set of dummy variables, and I allow XGBooost to determine which interactions are important. I weight each household according to the household weights provided in the ACS. Evaluated on a holdout sample, the three versions of the model correctly predict whether a household is above median income for 69.6%, 67.7%, and 66.7% of households, respectively. The imperfect precision indicates that there remain many unobservables that contribute to household wealth beyond a home PUMA and housing characteristics, which motivates using block group level income distributions to Bayesian-update the baseline predictions. I discuss this further below.

Given the model estimated on ACS data, I predict the probability that residents of each CoreLogic parcel are above median income based on the available parcel characteristics. When the parcel does not have an estimated home value—e.g., for a large apartment building with many units—the probability that residents are above median income depends on the age of the building, number of units, and PUMA.

The ACS PUMS data includes only the PUMA of the household, but in the parcels data I can observe their exact home location. Since cities are often quite income-segregated (Reardon and Bischoff, 2011), using the exact home location should substantially improve the precision of the estimates. To incorporate this information, I use the 2019 ACS 5-year block group income distributions to identify where above/below median income households live within the PUMA, then update the baseline probability that a parcel's residents are above median income using Bayes rule. For a tenant living in a parcel with characteristics x in block group g, I evaluate the probability they are above median income (H) as

$$P[H \mid g, x] = \frac{P[H \mid x]P[g \mid H]}{P[H \mid x]P[g \mid H] + P[L \mid x]P[g \mid L]}$$

where L denotes below median income and  $P[g \mid H]$  is the within-PUMA probability that a household of type H lives in block group g.

Finally, I use the mapping of device homes to parcels to assign a probability of being above median income to each device. 'Higher-income' devices are those for which the probability of being above median income, based on their home parcel, is at least 50%. Note that I do

not observe the actual income of the residents of any parcels, so I am unable to compute how much the Bayesian update step contributes to the precision of the estimates.

I repeat the same steps and data sources to predict whether or not a device is likely college-educated, which I use in Table 4 and Figure C.3. I define college-educated as having at least four years of college education.

#### A.2.4 Estimating a continuous measure of income

For some analyses, I use a continuous measure of household income (e.g., to define income group based tercile of the income distribution). I estimate income in levels following a similar procedure to the above. First, I estimate the relationship between household income (y) and housing characteristics  $(\mathbf{x})$  in the 2019 5-year ACS microdata. I use the same features as before and estimate the model using OLS. Using the estimated model, I compute the predicted income and the variance of the prediction error for each parcel in CoreLogic. Assuming the error is normally distributed, this provides an initial probability density function  $p(y | \mathbf{x}_i)$  for parcel characteristics  $\mathbf{x}_i$ .

Next, I update each estimate using the income distribution of the device's home block group. Intuitively, this will push a device's estimated income towards the average income of their block group, with larger changes for noisier estimates. The ACS tabulations include block group counts of households in 16 different bins of household income, ranging from 0.10,000 to > 250,000. Define  $h(b \mid \mathbf{x}_i, w_i)$  as the probability a household's income is in bin  $b \in B$  given parcel characteristics and home block group  $w_i$ . Under the assumption that the block group income distribution  $g(b \mid w_i)$  and the estimated distribution based on parcel characteristics  $p(y \mid \mathbf{x}_i)$  are independent,  $h(b \mid \mathbf{x}_i, w_i)$  is given by

$$h(b \mid \mathbf{x}_i, w_i) = \frac{g(b \mid w_i) \int_{\underline{b}}^{\underline{b}} p(y \mid \mathbf{x}_i) dy}{\sum_{b' \in B} \left[ g(b' \mid w_i) \int_{\underline{b'}}^{\overline{b'}} p(y \mid \mathbf{x}_i) dy \right]}$$

where  $\underline{b}$  and  $\overline{b}$  are the lower and upper household income bounds of bin b. The final income estimate is then a weighted average of bin midpoints with weights corresponding to the probability that the residents of a given parcel have a household income within that bin:

$$\widehat{y}_i = \sum_b \left[ \frac{(\overline{b} - \underline{b})}{2} h(b \mid \mathbf{x}_i, w_i) \right]$$
(A.1)

### A.3 Evaluating sample quality

#### A.3.1 Device sample

Evaluating the sample quality is complicated by the lack of any 'ground truth' about the demographics of devices in the data. To provide some sense of coverage, I compare the inferred home locations of device holders to the true population, using the 5-year 2019 ACS. Overall, the sample includes 6.1% of residents in the sample MSAs, although coverage is better in some MSAs (e.g., 9% in the Orlando-Kissimmee-Sanford MSA) than in others (e.g., 3.1% in the St. Louis MSA).

Next, I examine whether devices are disproportionately sampled from block groups with certain demographics. To do so, I divide block groups into within-MSA deciles by various characteristics, weighting by the population of each block group. Then, I compute the number of devices with homes in block groups corresponding to each decile.

Figure A.2 plots the fraction of devices from each decile for block group population density, median household income, median age, and the share of residents who identify as white (non-Hispanic). If the sample were perfectly uncorrelated with these characteristics, then 10% of devices would come from each decile (corresponding to the dashed horizontal line). Instead, on average, devices in the sample come from less dense, more middle-income, younger, and slightly less white block groups. One contributing factor is that Android devices are likely over-represented in the GPS data. While the data do not include the device type, anecdotally, vendors have said that iPhones are only 35-50% of their sample despite making up about 60% of US smartphones. Overall, the variation in coverage is similar to that documented by Couture et al. (2022) for a similar sample of GPS devices. Couture et al. (2022) also show that GPS data also matches well with other sources (e.g., travel surveys) when evaluating the distribution of trip distances and the rate of state-to-state migration.



Figure A.2: Device sampling by block group characteristics

*Notes:* This figure shows the sampling of devices from block groups relative to a uniform sampling. Dots above the line indicate that more than 10% of devices come from that block group. All block group characteristics are from the 2019 5-year ACS.

#### A.3.2 Establishments sample

To evaluate the coverage of establishments, I compare aggregate counts of establishments in each category to the corresponding counts from the County Business Patterns (CBP) for 2019. Figure A.4 plots the results for all SafeGraph establishments, split by counties with median household income in the top or bottom half of the sample of counties. SafeGraph has as many or more establishments of each category as the CBP, and the results are similar for counties of different income levels.

Figure A.3 plots the result when subsetting to SafeGraph establishments that are in the estimation sample for the model, where I restrict to establishments that have at least ten observed visits. Here, there are more establishments in the CBP than in SafeGraph. Some of the gap is also due to establishments that do not have their own footprint (e.g., stores in a large mall), which would appear in CBP but which I would classify as visits to the 'parent' establishment (i.e. the mall). But while there is a long tail of 'fringe' establishments in SafeGraph (especially for personal services), they make up a small share of visits -97.2% of visits by lower-income devices and 97.6% of visits by higher-income devices are to establishments that are in the final estimation sample.

### A.3.3 Visits sample

Table A.2 compares the distribution of daily trips to different categories of amenities in the GPS data and the 2017 National Household Travel Survey (NHTS). The NHTS surveys a single day of travel behavior for each respondent. I subset the NHTS data to the same set of MSAs as the GPS sample. The 'trip purposes' used for the NHTS do not perfectly align with the amenity categories that I define, especially for personal services and entertainment categories. Nonetheless, the number of daily trips looks reasonably similar in the two samples. In each sample, shopping trips are the most common, followed by entertainment (or 'social/recreation' in the NHTS), then restaurants, and finally services.

GPS dat	ta	NHTS d	ata
GPS category	# visits/day	NHTS purpose	# trips/day
Restaurants	0.190	Buy meal	0.252
Shops	0.575	Buy goods	0.436
Personal services	0.160	Buy services	0.070
Entertainment	0.248	Social/recreation	0.381

Table A.2: Visits by category: GPS data and the National Household Travel Survey (NHTS)

*Notes:* This table documents the number of daily trips in the GPS and National Household Travel Survey (NHTS) data for roughly comparable types of trips. The NHTS covers trips for a single day of activities for each individual. I use NHTS data from 2017 and restrict to the sample MSAs. For the GPS data, I compute the average number of visits to each category on days when the device is active.



Figure A.3: Estimation sample vs. County Business Patterns (CBP)

*Notes:* This figure compares counts of establishments in the 2019 County Business Patterns (CBP) to counts of establishments in the final estimation sample. Each plot is a binscatter, and gray bars represent 95% confidence intervals. I exclude establishments in the parks and malls subcategories as they are not covered by the CBP. Counties are split into above or below median income based on their household median income in the 2019 5-year ACS.



Figure A.4: SafeGraph vs. County Business Patterns (CBP)

*Notes:* This figure compares counts of establishments in the 2019 County Business Patterns (CBP) to counts of SafeGraph establishments. Each plot is a binscatter, and gray bars represent 95% confidence intervals. I exclude establishments in the parks and malls subcategories as they are not covered by the CBP. Counties are split into above or below median income based on their household median income in the 2019 5-year ACS.

# A.4 Trip chains

In modeling amenity consumption, I consider each choice to visit an amenity in isolation. In reality, individuals may combine many stays at different POIs into a single 'trip,' allowing them to amortize the cost of driving across multiple stops. This simplification could lead to overestimating the value of locations often combined with other stops.

To examine the frequency with which trip chains are observed in the GPS data, I combine stays at amenities into 'trips' by labeling a stay to be the start of a new trip of 1) the stay is at a POI that is not home or work and 2) the previous stay was at either home or work or the time between stays was two hours. Figure A.5 Panel (a) plots the distribution of number of POIs visits in a trip. 79% of trips consist of a single stop. For comparison, 61% of trips in the 2017 NHTS included multiple stops, although this is an imperfect comparison as the NHTS 'tours' include non-POI stops that are not part of the GPS trip chains (e.g., at school or a friend's house). Panel (b) plots the distribution again, breaking out by device coverage quartiles. I define coverage at the device-week level as the sum of all time observed in the given week – on average, the top quartile of device-weeks is observed for about 83% of all minutes of the week, while the bottom quartile is observed for about 12%. For devices in the top quartile of coverage, 71% of chains consist of a single stop, compared to 84% of trips in the bottom quartile. Finally, Panel (c) shows that the distribution of chain length is similar for higher and lower-income residents.

# A.5 Identifying 'displaced' residents in Infutor

To identify where potentially displaced long-time residents move, I use data on individual address histories from Infutor, which covers most of the US adult population and has been shown to be broadly representative (Phillips, 2020). I define an individual as potentially displaced from a Census block group if they meet the following criteria:

- 1. They have lived in the block group for at least five years as of 2015, and are between 18 and 75 years old that year
- 2. The block group's median rent increases by at least 25% between the 2015 5-year ACS and 2019 5-year ACS
- 3. The block group is 'high-amenity' by 2019, which I define as the top quartile of the NAQI distribution for their MSA
- 4. They move to a neighborhood in their MSA with cheaper median rents. (The restriction to within-MSA moves is because NAQI values are only available for the 30 most populous MSAs)

In total, I identify 2,606 unique block groups that meet this definition of gentrifying out of 57,661 total block groups in the sample MSAs that have non-null median rents in 2015 and 2019. From these block groups, I identify 109,326 moves out by long-time residents, which represents 16.6% of the population in these block groups who had lived there at least five years and were aged 18-75 in 2015.



Figure A.5: Distribution of # stops in trip chain

*Notes:* These figures document the distribution of the number of stops at different establishments within a single trip chain, where a new trip begins if a device stops at home or work or it has been at least two hours since the device left the previous establishment.

# **B** Model appendix

# **B.1** Estimation details

**Estimation of lowest level.** For each subcategory, I use data on all visits to establishments within this subcategory to estimate the establishment-level taste parameters and the disutilities to driving time from home and work.<sup>35</sup> I use a fixed choice set consisting of all establishments in a subcategory for which I observe at least five visits in the data. For many subcategories, the choice sets are large; for example, there are over 12,000 full-service restaurants in Los Angeles.

 $<sup>^{35}{\</sup>rm There}$  are a few cases of driving times being unreasonable long for within-MSA trips, likely due to errors in the router, so I top-code driving times at 90 minutes and include an indicator for whether a driving time was over 90 minutes.

For estimation, I use a mini-batch gradient descent algorithm built in PyTorch and run it on a virtual machine with 16 CPUs, 104GB of memory, and 1 Nvidia T4 GPU with 12GB of GPU memory. 'Mini-batch' refers to estimation by iterating over smaller batches of data in this case, with 50,000 choices in each—so that all the data does not need to be held in memory at once (with large choice sets, even 50,000 choices becomes prohibitively large to store within a GPU's memory). I use an Adam optimizer with an initial learning rate of 1e-2 and a convergence tolerance of 5e-7. Adam uses only the current and historical gradients (no Hessian) and adjusts the learning rates for each parameter separately. I weight the loss function by a given device's sample weight. I represent fixed effects for C different establishments in a choice set using a  $1 \times C$  vector of embeddings, which is functionally equivalent—but far more memory efficient—than using one-hot encoding ('dummy variables') way to represent high-dimensional fixed effects.

There are a few cases where an MSA will have zero 'gambling' establishments, and so that subcategory is excluded from estimation.

**Estimation of upper levels of model.** To link the lowest level to the upper levels, I compute the inclusive values for each subcategory as

$$DTE_{itl}^{k} = \sigma_{l}^{k} \log \left( \sum_{l \in m} \exp \left[ (\gamma_{j}^{k} - \kappa_{t}^{k} \mathbf{d}_{ij}) / \sigma_{l}^{k} \right] \right)$$
(B.1)

where  $\sigma_l^k$  is the scale parameter of the  $\varepsilon_{ijt}$  for establishments within subcategory *l*. Similarly, the inclusive values that link the middle level (choice of subcategory) to the upper level (choice of category) are given by

$$IV_{itm}^{k} = \log\left(\sum_{l \in m} \exp\left[(\gamma_{l}^{k} + \beta^{k} IV_{itl}^{k})/\rho_{m}^{k}\right]\right)$$
(B.2)

To estimate the upper levels of the model, I again use PyTorch with an Adam optimizer for estimation with an initial learning rate of 1e-3 and convergence tolerance of 1e-7. Since the data and number of parameters are more reasonable, I do not batch the data; instead, I compute the loss and update the parameters using the full data for each 'step' in the optimization. To help avoid finding only local maxima—the objective function of a nested logit is not globally convex—I use a cosine annealing learning rate scheduler that increases the learning rate every so often to try to jump out of any local maxima. In testing, different initializations of the parameters led to similar final estimates.

## **B.2** Parameter estimates

Figure B.1 plots the disutility of a minute of driving from home or work by time of week, normalized so that the disutility from home during the weekday evenings is -1. On weekdays, both higher and lower-income devices become more sensitive to driving further from home and less sensitive to driving further from work as the day progresses. On weekends, individuals effectively only care about driving time from home. Table B.1 documents the category-time fixed effects  $(\gamma_{mt})$  and the degree of independence for each category  $(\rho_m)$ . Each parameter is the population-weighted average across all MSAs. Subcategories within shops exhibit the most independence from each other, while personal service subcategories exhibit the least. The value of different categories varies meaningfully by time of week. For example, individuals place relatively more value on restaurants in the afternoons and evenings and more value on shops during the mornings.





*Notes:* This figure illustrates how the disutility of driving time varies by time of day. Values are normalized relative to the disutility of driving from home in the evenings. Each coefficient is averaged across all MSAs, weighting by their population. Grey bars represent 95% confidence intervals.

Time of week	Category	Н	L
Nest independence $(\rho_m)$			
	Restaurants	0.50	0.45
A 11	Shops	0.73	0.72
All	Personal Services	0.35	0.45
	Entertainment	0.57	0.49
Category intercepts $(\gamma_{mt})$			
	Restaurants	-3.03	-2.89
Weekday mornings	Shops	-1.62	-1.69
	Personal Services	-1.81	-1.81
	Entertainment	-3.22	-3.01
	Restaurants	-1.60	-1.58
Wookday afternoons	Shops	-1.53	-1.53
Weekuay anermoons	Personal Services	-2.18	-2.10
	Entertainment	-0.87	-0.97
	Restaurants	-1.64	-1.60
Wookday avaning	Shops	-2.46	-2.39
weekday evenings	Personal Services	-0.82	-0.90
	Entertainment	-1.57	-1.54
	Restaurants	-1.94	-1.75
Weekend morning	Shops	-1.28	-1.18
weekend mornings	Personal Services	-2.20	-1.93
	Entertainment	-2.33	-2.13
	Restaurants	-1.72	-1.55
Weekend offernaans	Shops	-2.37	-2.09
weekend alternoons	Personal Services	-2.30	-2.11
	Entertainment	-1.55	-1.50
	Restaurants	-1.77	-1.61
Weekend evening	Shops	-2.39	-2.34
weekend evenings	Personal Services	-1.39	-1.39
	Entertainment	-1.89	-1.69

*Notes:* This table documents the category intercept and nest independence parameters for the upper-level model, averaged across all MSAs (weighted by population).

# C Supplementary exhibits



Figure C.1: Additional measures amenities access

(a) Amenities within 10 minutes (w/ log pop. desnity controls)

*Notes:* The first panel replicates Figure 1—where the outcome is the number of POIs within a 10-mile drive—adding controls for log population density of the block group from the 2019 5-year ACS. The second panel plots the number of POIs within 1 mile 'as the crow flies.' Each point is the coefficient from a regression of the inverse hyperbolic sine of the outcome on indicators for income quartile using device-quarter level of data, with MSA fixed effects. (I use the inverse hyperbolic sine to handle zeros, although I refer to it as 'log' for brevity). Gray bars represent 95% confidence intervals, with standard errors clustered by home or work block group.



Figure C.2: Neighborhood amenity consumption by race/ethnicity

*Notes:* This figure plots the number of visits to amenities of each category, split by White devices (WD) and non-White devices (NWDs). A device is labeled as a WD if its home Census block was over 50% non-Hispanic White in the 2010 Census. Each point is the coefficient from a regression of the inverse hyperbolic sine of the outcome on indicators for income quartile interacted with whether WD. I use the inverse hyperbolic sine to handle zeros, although I refer to it as 'log' for brevity. The data are at the device-quarter level of data. Income quartile cutoffs are based on the distribution of incomes in the 30 MSAs in the sample. Gray bars represent 95% confidence intervals, with standard errors clustered by device.



Figure C.3: Neighborhood amenity consumption by whether college-educated

*Notes:* This figure plots the number of visits to amenities of each category, split by whether a device is college-educated. Predicting education follows the same steps as predicting above/below median income (i.e., using parcel characteristics, with a Bayesian update step based on the demographics of their home block group. See Appendix Section A.2.3 for details.). Each point is the coefficient from a regression of the inverse hyperbolic sine of the outcome on indicators for income quartile interacted with whether college-educated. I use the inverse hyperbolic sine to handle zeros, although I refer to it as 'log' for brevity. The data are at the device-quarter level of data. Income quartile cutoffs are based on the distribution of incomes in the 30 MSAs in the sample. Gray bars represent 95% confidence intervals, with standard errors clustered by device.



## Figure C.4: Visits to major brands by income quartile

*Notes:* This figure documents the share of visitors by income quartile to major brands in four different subcategories. Income quartile cutoffs are based on the distribution of incomes in the 30 MSAs in the sample.



Figure C.5: Relative value of restaurant chars. (with tract fixed effects)

(a) Cuisines

Notes: This figure documents the relationship between the restaurants' estimate value for each income group and different observables. Each point corresponds to a coefficient of the establishment-level fixed effects  $(\gamma_j^k)$  on characteristics of the establishment with controls for the tract of the establishment. Gray bars represent 95% confidence intervals.

	Std. dev. of $\varepsilon$		Percent	of visits
	Н	L	Н	L
Restaurants				
Full-service	9.39	9.94	11.28	10.98
Limited-service	8.76	9.03	3.04	3.45
Cafes; snacks	8.89	9.24	2.01	1.70
Drinking places	9.77	10.01	1.49	1.48
Shops				
Malls	8.09	8.41	31.11	29.50
General merch., warehouse clubs	6.40	6.66	3.78	5.36
Clothing, shoes, jewelry, leather goods	9.37	9.80	3.85	3.70
Groceries, beer/wine/liquor stores	6.55	6.92	3.41	3.83
Gas stations, convenience	10.13	10.51	2.07	3.26
Department	6.62	6.95	2.07	2.02
Building materials, gardening	6.79	7.18	1.76	2.14
Furniture, appliances, electronics	9.12	9.40	1.45	1.34
Automobile dealers	11.72	11.82	1.34	1.37
Sporting goods, hobby, music	9.10	9.50	1.18	1.06
Pharmacies	6.98	7.34	1.04	1.21
Beauty, glasses, personal care	8.49	8.81	1.03	1.00
Books, office supplies	9.35	10.04	0.60	0.55
Personal services				
Hospitals, health clinics	11.93	11.38	3.73	4.42
Salons, barbers	8.41	8.69	2.44	2.39
Religious organizations	8.87	9.22	1.60	1.57
Banks	7.64	7.91	1.28	1.52
Dentists	8.55	9.13	0.84	0.78
Car maintenance	9.85	9.61	0.28	0.36
Drycleaning, laundry	6.74	6.32	0.09	0.11
Entertainment				
Parks	8.37	8.57	7.75	7.30
Fitness centers	8.01	7.59	2.87	2.17
Golf courses	8.14	9.57	2.70	1.64
Movie theaters	8.36	8.99	1.01	0.90
Performing arts and spectator sports	18.26	19.59	1.02	0.87
Other amusement, recreation	10.34	10.45	0.70	0.65
Libraries, museums, zoos, gardens	8.57	9.10	0.49	0.44
Gambling	14.11	13.61	0.37	0.58
Amusement parks	10.75	11.89	0.35	0.35

Table C.1: Subcategory  $\varepsilon$  variance and share of visits

*Notes:* For each subcategory, this table documents the share of each income group's visits to establishments in that subcategory and the standard deviation of the idiosyncratic component of utility across establishments (conditional on subcategory choice). The latter is estimated as part of the lower level estimation steps and is units of minutes of driving time. All statistics are averaged across MSAs, weighting by the MSA population.

# Figure C.6: NAQI values

#### (a) Houston



*Notes:* This figure illustrates the estimated block group level NAQI values for different MSAs, zoomed in slightly to focus on the urban core. NAQI values for each group correspond to minutes of weekly driving time relative to the median neighborhood (for that income group).

Table C.2: Results by MSA

			NA	QI	С	ounterfac	tuals
	Corre	elation	p75-	-p25	Ta	ilor	Displace
	POIs	NAQI.	Н	L	Н	L	All
Los Angeles-Long Beach-Anaheim	0.73	0.98	46.38	44.32	6.94	-4.20	-29.56
Chicago-Naperville-Elgin	0.74	0.98	86.45	78.91	8.65	-6.86	-34.36
Dallas-Fort Worth-Arlington	0.77	0.99	69.82	73.81	5.03	-6.14	-32.66
Houston-The Woodlands-Sugar Land	0.78	0.98	89.97	91.38	5.61	-4.61	-40.23
Washington-Arlington-Alexandria	0.73	0.99	101.93	107.38	9.84	-8.39	-40.30
Miami-Fort Lauderdale-West Palm Beach	0.74	0.99	74.85	78.53	9.58	-5.98	-26.93
Philadelphia-Camden-Wilmington	0.72	0.98	80.40	84.05	6.86	-4.95	-32.37
Atlanta-Sandy Springs-Roswell	0.78	0.99	119.96	101.49	7.36	-5.83	-52.10
Boston-Cambridge-Newton	0.68	0.98	107.89	135.34	11.47	-10.61	-67.93
Phoenix-Mesa-Scottsdale	0.76	0.99	71.06	55.89	8.77	-5.26	-34.51
San Francisco-Oakland-Hayward	0.68	0.99	68.12	75.46	11.44	-6.04	-37.86
Riverside-San Bernardino-Ontario	0.71	0.97	132.30	109.27	11.27	-4.81	-49.07
Detroit-Warren-Dearborn	0.71	0.98	64.58	58.51	8.86	-6.26	-37.36
Seattle-Tacoma-Bellevue	0.67	0.97	82.60	83.35	12.18	-11.57	-47.70
Minneapolis-St. Paul-Bloomington	0.72	0.98	72.62	85.32	6.33	-14.98	-48.47
Tampa-St. Petersburg-Clearwater	0.71	0.98	83.35	80.22	6.03	-8.94	-40.45
Denver-Aurora-Lakewood	0.72	0.99	50.68	52.33	5.94	-4.62	-39.22
St. Louis	0.65	0.99	115.95	106.64	8.84	-8.91	-34.77
Baltimore-Columbia-Towson	0.66	0.98	86.29	101.34	10.45	-15.46	-54.19
Charlotte-Concord-Gastonia	0.71	0.99	109.42	120.74	5.80	-7.52	-42.46
Orlando-Kissimmee-Sanford	0.74	0.99	98.48	129.37	5.02	-7.62	-81.69
San Antonio-New Braunfels	0.73	0.98	60.16	64.91	6.58	-6.70	-46.15
Portland-Vancouver-Hillsboro	0.67	0.99	72.08	74.18	9.32	-9.35	-40.00
Pittsburgh	0.69	0.99	106.15	103.79	8.50	-5.58	-38.63
Sacramento–Roseville–Arden-Arcade	0.68	0.99	73.44	93.79	8.77	-11.43	-39.82
Las Vegas-Henderson-Paradise	0.66	0.98	61.01	52.42	11.65	-8.30	-57.33
Cincinnati	0.72	0.98	92.71	85.48	5.43	-6.89	-55.59
Kansas City	0.62	0.98	81.68	79.19	11.92	-10.02	-32.47
Austin-Round Rock	0.71	0.97	100.63	91.58	7.25	-7.45	-42.31
Columbus	0.67	0.98	83.36	107.59	7.70	-7.98	-32.59

*Notes:* This table documents the primary set of results separately for each MSA in the sample. The first two columns document the correlation in establishment-level preferences subcategories (weighted by the total number of visits in each subcategory) and NAQI values. The second two document the interquartile range for the NAQI values, which are in units of minutes of driving time per week. The tailoring counterfactual document the effects (in minutes/week of value) of tailoring 25% of establishments to the preferences of higher-income residents. The final column documents the average change in NAQI value for long-time residents of gentrifying neighborhoods that moved to a cheaper neighborhood between 2015-2019.