Housing Affordability and the Mobility of Renters*

Nathaniel Baum-Snow†, Abdollah Farhoodi‡, Lu Han§

September 13, 2025

Abstract

This paper examines the welfare costs of rent growth and the prospects for reducing these costs. Using panel data on the near universe of renters in Texas 2011-2019, we estimate migration costs between neighborhoods jointly with the demand system for neighborhoods and commuting zones. A structural dynamic model of demand for cities and neighborhoods with endogenous mobility underlies these estimates. Using estimated model parameters, we quantify how migration to substitute cities and neighborhoods has mitigated the welfare cost of rent growth by household type. A reduced-form analysis of the impacts of housing supply shocks on renters' mobility and rental payments corroborates this model-based decomposition.

*We thank seminar audiences and in particular Victor Aguirregabiria, Sophie Calder-Wang, Gilles Duranton, Greg Howard, Pierre Mabille, Kyle Mangum, Jordi Jofre Monseny, Yuhei Miyauchi, Matthew Osborne, Diego Puga, Marlon Seror, Chris Severen, Bryan Stuart, Nicholas Trachter, and Yanos Zyllerberg for many helpful comments. James Macek provided superb research assistance.

†Professor, University of Toronto, Rotman School of Management

‡Assistant Professor, University of Toronto, Department of Economics

§Professor, University of Wisconsin-Madison, Wisconsin School of Business

1 Introduction

Between 2014 and 2022, the Zillow rent index for the US increased by 61%, which was more than twice the rate of general price inflation. The extent to which this rental price growth burdens incumbent renters depends not only on rent changes in their neighborhoods, but also on the cost of moving, rents in alternative neighborhoods, and the overall value of moving relative to staying. While the mobility option plays a potentially central role in mitigating the welfare burden of rent growth, heretofore there is little systematic evidence on the value of this option. Understanding the extent to which household mobility can facilitate housing affordability is central to designing policies that mitigate the welfare burdens of rent growth. Quantification of the underlying mechanisms driving renters' mobility responses to housing market shocks delivers valuable insights into the consequences of various supply-side policies that influence rents.

This paper systematically investigates the welfare consequences of rising rental prices for renters. We build an individual-level panel data set of residential locations for the near universe of renters in Texas by marrying migration histories from the Infutor marketing data set with housing tenure information in CoreLogic. We develop a dynamic structural model of location choice and endogenous migration with parameters that can be feasibly estimated with this data set. Using parameter estimates from the model, we quantify the welfare costs of rent growth between 2010 to 2019 for incumbent renters in different types of neighborhoods. In particular, we quantify the extent to which migration has allowed incumbent renters to mitigate the welfare burden of rent growth in their neighborhoods. Essential to this quantification are credible estimates of two sets of parameters. First, we recover estimates of parameters governing migration costs across neighborhoods within and between commuting zones. Second, we recover estimates of utility parameters by accounting for the existing dynamic location choice decision and governing the demand system for commuting zones and neighborhoods in a flexible dynamic framework. Associated population supply elasticities are allowed to differ by commuting zone size and location, individual characteristics, and origin neighborhood 2010 characteristics.

In the model, individuals face a move-stay decision each year. This is the top nest of a nested logit. Conditional on moving, we explore two model specifications. In the first, substitution patterns across neighborhoods are governed by a flexible function of neighborhood and individual characteristics in a logit demand system. In the second specification, choices are nested into commuting zones and neighborhoods within commuting zone. In both specifications, each neighborhood delivers a flow utility that depends on the market rent index, a neighborhood fixed effect, CBSA-time fixed effects, an amenity shock that is common to all residents, and an idiosyncratic shock. Similar to Diamond et al. (2018), migration costs and flow utilities depend on the number of years in the residence and number of years in the

commuting zone. Over time, people can establish stronger ties to their home and city. The magnitude of local ties determines the extent to which housing affordability can be mitigated through providing lower rent housing options in far away neighborhoods or in other commuting zones. Moreover, these tenure effects allow the model to fit the "root-t" cumulative out-migration rate in tenure in the home observed by Howard and Shao (2023). Migration costs also depend on distance migrated and whether migration is within or between commuting zones. Additionally, individuals receive a flow utility shock associated with staying in their home. The parameter governing the variance of this shock captures the likelihood of lease non-renewal by the landlord, changes in family structure, or changes in employment situation that influence individuals' housing choices independent of their neighborhood, and can be thought of as a general preference for moving. Time-varying state variables in the model are location, year, residence tenure, and commuting zone tenure.

As full information maximum likelihood estimation using individual level panel data is computationally infeasible, we employ two complementary strategies to estimate the model. Our primary strategy uses a neural net machine learning algorithm in a first step to predict movers' migration decisions as a function of state variables. This allows for calculation of an individual-specific summary measure of the value of moving to alternative homes along with decomposition into variable migration cost parameters in a second step. A third step of estimation uses the move-stay decision to identify the value of staying in a neighborhood plus the fixed cost of moving and a parameter governing the preference for moving. In the fourth and final step, we separately identify the fixed cost from parameters governing flow utilities. We recover the fixed moving cost (for moves of 0 distance) essentially by using the predicted rate of moves across homes within neighborhoods.

As robustness, we carry out a parallel analysis using data aggregated to the migration flow-year by residence and commuting zone tenure level in a nested logit framework. Parameters governing variable moving costs and substitution elasticities are estimated in a first step using an iterative high dimensional fixed effects Poisson pseudo maximum likelihood (PPML) estimator. Destination fixed effects in this migration gravity estimation equation identify the conditional values of each location. Both the neural net and PPML estimator implement the Hotz and Miller (1993) inversion without requiring explicit smoothing like that employed in Bayer et al. (2016) while accommodating Dingel and Tintelnot (2020)'s guidance on the handling of sparse migration flow matrices. With aggregate data, the second step of estimation uses the fact that the move-stay decision in period t relative to $t+1$ depends on the neighborhood flow utility, including the zip code level Zillow rent index, and outside opportunities, as captured by a composite value of moving. The renewal process attribute of the model allows us to difference out unknown continuation values, representing the dynamic move-stay decision in a simple regression equation with an expectation error.

We instrument for rents using a novel shift-share instrumental variable: the interaction of neighborhood-level housing supply elasticity estimated in Baum-Snow and Han (2024) and the overall trend in rental prices. We demonstrate that this variable predicts two reasures of neighborhood rent conditional on a rich set of fixed effects. However, it does not predict changes in measures of neighborhood amenities, as is required for identification of our demand model.

Dynamic model parameter estimates indicate a number of revealing patterns. First, moving costs are much higher between than within commuting zones, and the between commuting zone moving cost increases markedly with commuting zone tenure. Second, younger individuals have lower moving costs and lower income individuals have higher moving costs. Third, larger population commuting zones have higher moving costs within them but lower moving costs to them. Fourth, neighborhoods in individuals' own commuting zones are poorer demand substitutes in larger population commuting zones, though substitution elasticities between neighborhoods in other commuting zones do not depend on their populations. Individuals gain utility from remaining longer in the same commuting zone. Consistent with the idea that positive neighborhood amenity shocks are capitalized into rents, OLS estimates of the marginal utility of log rent are biased upwards (toward 0). Consistent with the idea that positive neighborhood amenity shocks are serially correlated at an annual frequency, OLS estimates of demand substitution elasticities across neighborhoods are biased downwards (toward 0).

Estimation of a static version of the model by PPML delivers patterns of moving costs and marginal utilities that are consistent with dynamic model estimates. This is true even though the marginal utilities of rents are estimated through migration destination choice in the static model rather than the move-stay decision in the dynamic model.

Using the model, we account for the extent to which growth in neighborhood rents has been burdensome for residents. In commuting zones without alternative low rent neighborhoods available, we see that growing rents has large negative impacts on welfare. With our estimates, counterfactual analysis can be used to evaluate the efficacy of various policies that enhance the availability of affordable housing. Considering the effects on welfare through associated changes in housing rents and migration patterns helps to evaluate the relative welfare benefits of demand-side rental subsidies (as in Section 8 housing vouchers) versus supply-side construction subsidies (as in the Low Income Housing Tax Credit) or relaxed land use planning and inclusionary zoning policies. Moreover, our estimated model allows for evaluation of the extent to which targeting such policies to specific neighborhoods is welfare-enhancing. Counterfactual analysis using our estimated model also allows us to assess the efficacy of place based policies that enhance the desirability of living in low cost housing markets or of reducing migration costs.

We contribute to the extensive literature on housing affordability Molloy (2020). This paper distinguishes itself by focusing on renters as the most vulnerable group to housing shocks, an aspect that has received less attention in existing studies (Diamond et al., 2019). As rents have risen less quickly than home prices in supply constrained markets (Molloy et al., 2022), a focus on the welfare burdens of changes in rental affordability is also informative for higher income populations who may become newly dissuaded from purchasing a home. Additionally, our study aligns with Gyourko et al. (2013) on the migration of low-income households from high-price locations and Ganong and Shoag (2017) on the relatively low long-distance migration rates of low-skilled workers, even those that are exposed to rising housing costs. Finally, we draw inspiration from a large literature on dynamic location choice models to estimate our structural model (Bayer et al., 2016; Caliendo et al., 2019; Couture and Handbury, 2020; Schubert, 2021; Howard and Shao, 2023; Almagro and Domınguez-Iino, 2025).

2 Data and Descriptive Evidence

2.1 Data

We use Infutor marketing data, CoreLogic's compendium of residential housing assessment data, Zillow Observed Rent Index (ZORI) data, American Community Survey (ACS) micro data (Ruggles et al., 2020), and 5 year neighborhood aggregates of ACS data for this analysis. We focus our analysis on the state of Texas for the 2010 to 2019 period. Texas is a convenient benchmark as it has both a large and heterogeneous set of commuting zones and neighborhoods and little rental market regulation. In addition, the fact that it was rapidly growing during this time period helps with implementing the identification strategy we lay out below. Almost all Texas neighborhoods experienced rental housing demand growth during our study period.

The ZORI is our main source of information about market rents. This is a "repeat rent" index of asking rents constructed using data from all rental listings on the Multiple Listing Service plus some additional proprietary listing data collected by Zillow. It closely follows the methodology in Ambrose et al. (2015), with the addition of weights to reflect the age and units in the structure composition of the rental stock in each year. We note that actual paid rents for new units may lag this index by up to several years as lease renewals are not included in the underlying data.

As the ZORI data comes at the zip code level of geography, we are constrained to use neighborhoods that are geographically defined as zip code aggregates. The neighborhood is the core spatial unit in our analysis. We employ an algorithm that combines adjacent

zip codes into neighborhoods with populations with a minimum target of 6,000 residents in 2010. Within each metropolitan area (CBSA), the algorithm combines adjacent small municipalities into neighborhoods, assigns each medium sized municipality to be its own neighborhood, and partitions large municipalities into multiple neighborhoods. As zip codes vary greatly in size, with many that are very large, we end up with a median population of 20,000 in 1,029 Texas neighborhoods. Of these, 435 are in the large city size category that includes the Houston and Dallas-Fort Worth CBSAs (total 2010 population 14.9 million), 180 are in the medium size category of Austin, San Antonio and El Paso (total population 5.7 million), and the remainder are in the small/rural size category (total population 8.7 million). Using the 2000 census and subsequent waves of 5-year ACS micro data aggregated to 2020 Zip Code Tabulation Areas (ZCTAs), we construct information about the housing market and demographic characteristics for each neighborhood over time. Such measures include distributions of gross rents, rent to income ratios, age, education, and income. Commuting zones conform to 2020 CBSA definitions and the state remainder.

We spatially aggregate the Zillow rent index across neighborhood zip codes in each year using the number of rental units from CoreLogic as weights. In any given year, we do not observe rents for between 12 and 22 neighborhoods.[1] Missing neighborhoods have very few rental units. We use a version of the ZORI that covers the 2010-2019 study period.

We utilize the comprehensive Infutor marketing data set as our primary data source to investigate migration patterns. This data set incorporates information on the residential addresses of US households, including their ten most recent addresses. It covers most US residents back to at least 2006, with dates and addresses of residential moves. Infutor compiles these data from publicly available sources, including change of address requests from the Post Office, credit card records, property deeds, public education records, and state criminal record databases. Although primarily marketed for customer relationship management purposes in the private sector, the Infutor data have proven valuable in other recent academic work, notably including Diamond et al. (2019).

Our study requires us to identify renters in the Infutor data. To do so, we match each individual's last home address in Infutor to that in the CoreLogic home assessment data. The resulting match rate is 96%, including fuzzy matches using addresses of housing complexes that are missing unit numbers or that match imperfectly in other ways. CoreLogic includes a variable that indicates whether the resident of a housing unit owns or rents, based on whether the property tax bill is sent to a different address. Based on this variable, the last observed address in Infutor, and the fuzziness of the Infutor-CoreLogic address match, we assign a probability to each individual indicating whether they are a renter or an owner. Successful exact address matches (about 60%) adopt the CoreLogic assigned probability.

[1] In most years, 952 Texas zip codes have a rent index value calculated.

Those with a matched building but not unit number take on the inferred fraction of building residents that are renters. Those with a matched street but not number adopt the fraction on the street in that zip code that are renters. We keep all individuals in the sample with a probability of renting of at least 0.7 in their last address observed in Infutor. The result is a data set of 7.7 million renters in Texas for the 2010-2020 period on a base of 25.1 million residents in 2010.[2]

Using the Infutor data, we build a transition matrix out of each neighborhood n in each commuting zone l in each year t by home and commuting zone tenures ι^r and ι^l. We incorporate observed household heterogeneity into the analysis using origin ZCTA characteristics, including income and age as measured in the 2008-2012 ACS, and age, income group, and presence of children in the household, as reported by Infutor.[3] Movers may choose a different home in the same neighborhood, any other neighborhood in Texas, or the "outside option" outside Texas. Greater flows in this matrix indicate some combination of lower pairwise inter-regional migration costs and a greater value associated with the destination location in year t. Table A2 summarizes the resulting calculated renter flows. We measure home and CBSA tenure using Infutor's report of when it begins tracking each individual. Those entering the data later are assigned random tenures drawn from the distributions for their neighborhood on entry into the data.

To corroborate Infutor's tracking of the distribution of origin-destination migration pairs conditional on moving, we compare this flows matrix to that in the ACS data for each year 2012-2019.[4] The ACS reports the Migration PUMA, of which there are 71 in Texas, in which each individual lived one year prior to being surveyed. To make comparisons, we aggregate Infutor flows to be between these 71 spatial units. To compare the two data sets, we calculate the fraction of within Texas out-migrants from each migration PUMA in each year migrating to each destination. This distribution of conditional choice probabilities (CCPs) in Infutor contains the key empirical moments used in estimation to recover moving costs and renewal values associated with each neighborhood.

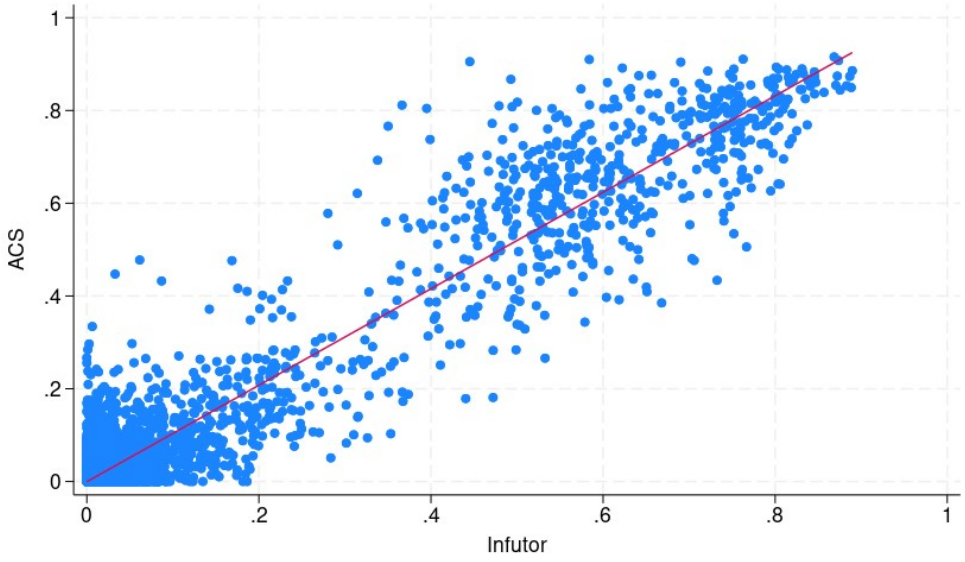
We do not expect these pairwise flows to perfectly align for two reasons. First, the ACS is only a 1% random sample of the population, making its migration flow measures noisier than Infutor's. Second, the Infutor data reports actual moves whereas the ACS data only reports residential locations one year apart. Nevertheless, we see a strong relationship between migration flow measures in these two data sets. Among the 40,328 neighborhood pairs in each year 2012-2019, the correlation is 0.9 and is robust to year, origin, and destination fixed

[2] 35 percent of housing units in Texas were renter-occupied in 2010. Given the total population growth in Texas to 29.1 million in 2020 and our count of 18.0 million owners in the Infutor data, we conclude that the Infutor data misses some renters and we have likely mis-classified a small share of renters as owners.

[3] We 2022 or 2014 demographic information from Infutor for about one-half of individuals. This information is inferred in part use the census demographic information for the individual's residential neighborhood.

[4] We begin in 2012 as this is the first year that 2010 census geographies are used in the ACS.

Figure 1: Comparison of Renter Migration Flows in the ACS and Infutor



Note: Points show conditional choice probabilities of migration destinations for renters in the ACS and Infutor data sets in each year 2012-2019. Each observation is a Migration PUMA pair for a given year.

effects. Given measurement error in ACS flows, we expect this correlation to be less than 1. Figure 1 presents a scatterplot comparing these CCPs in each data set.

We note that Infutor is better at tracking migration flows than at measuring population stocks. Corroboration with ACS aggregate data reveals that Infutor typically undercounts long-term stayers. In estimation, it will be important to get the move probability right out of each neighborhood, meaning we need accurant counts of stayers. As such, we weight each stayer constructed from the Infutor data by the ratio $\frac{Pop^{ACS}-Pop^{Infutor}_{mover}}{Pop^{Infutor}_{stayer}}$ separately by neighborhood and year, only using renters. This ensures that total population counts match between our estimation sample and the more representative ACS data.[5]

2.2 Patterns of Rents and Affordability

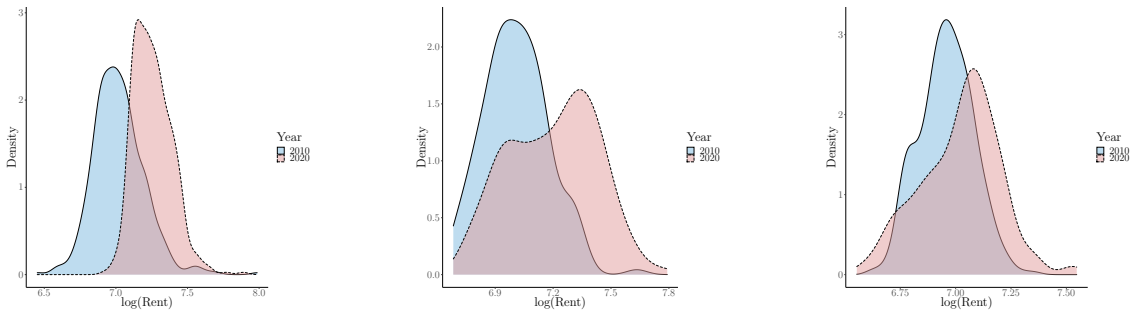
Next, we document some stylized facts that guide key modeling decisions used in the following section. Descriptive patterns documented in this section also point to some of the qualitative conclusions that emerge after model estimation. In particular, our evidence of rising rents in almost all neighborhoods of large CBSAs married with continued positive net migration of renters to these metros indicates rising amenity values for these locations. Moreover, the

[5]We use 5 year aggregate ACS data for this adjustment, indexed to the middle year. As it smooths over time and uses a 5% sample, this procedure introduces both idiosyncratic noise in the number of stayers and a more sluggish response of move probabilities to neighborhood shocks. Both of these limitations are addressed in estimation.

high mobility rates of renters within but not between metros will be rationalized in the model as a combination of differences in moving costs and neighborhood substitution elasticities. We then build up to our identification strategy for isolating variation in neighborhood rents, which is key to credible estimation of utility function parameters in the model.

Figure 2 provides insight by examining large, medium and small size CBSAs. Over 10 years, the rent distribution shifts to the right across all types of CBSAs with a marked increase in the rent level, particularly so for large and medium size CBSAs, but makes a small change in rent dispersion across markets within each type. A striking feature of the data is that rental prices in the more affordable neighborhoods within large CBSAs have risen significantly, and large CBSAs have very few low rent neighborhoods left after 10 years. Figure 2-a shows this with its kernel density plot of the rent distribution in Dallas and Houston in 2010 and 2020: neighborhoods with rent below the median in 2010 no longer existed in 2020.

Figure 2: Shifts in Rent Distributions Across Neighborhoods Over Time



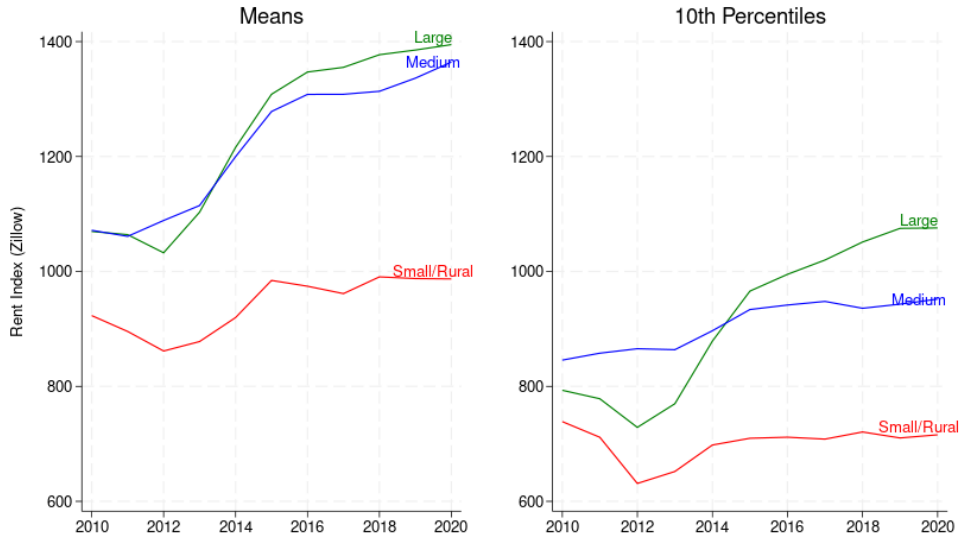
(a) Large CBSAs (b) Medium CBSAs (c) Small CBSAs

Note: Large CBSAs include neighborhoods in Dallas and Houston. Medium CBSAs include neighborhoods Austin, El Paso and San Antonio. Small CBSAs include all remaining neighborhoods.

Looking at variation in rental prices over time, Figure 3 presents patterns in Zillow rent indices and average gross market rents experienced by the average (left panel) and 10th percentile (right panel) renter across rural areas/small CBSAs, medium sized CBSAs, and large CBSAs since 2010. Medium and large CBSAs have experienced marked increases in rents since 2012. However, small CBSAs and rural areas remain relatively affordable. Moreover, while 10th percentile neighborhoods in medium sized CBSAs are still affordable relative to 2010, even 10th percentile neighborhoods in large CBSAs experienced rapid rent growth since 2012. One object of our analysis below is to determine the welfare costs of this rise in unaffordability in the lowest rent neighborhoods.

While rents rose in the 2010-2019 period, renters' incomes rose approximately equally quickly. Increases in Texas rent to income ratios of 0.26 to 0.31 between 2000 and 2007

Figure 3: Patterns of Rents 2000-2020



Note: Lines depict the Zillow rent index (ZORI) averaged across zip codes in each indicated CBSA size category weighted by the number of rental units in each year as measured in the ACS. The Zillow rent index is calculated relative to 2020 observed gross rents. All rents are in current (undeflated) dollars. The left panel shows weighted averages across neighborhoods. The right panel shows indexed and mean gross rent for neighborhoods in the (weighted) 10th percentile of each rent distribution.

stabilized at the higher level till 2019. For 2007 to 2013, this stability went along with a small increase in the fraction of households renting from 0.34 to 0.37, with stability till 2019. However, this overall stability in rent to income ratios masks some heterogeneity across CBSAs. As seen in Figure A1, rent to income ratios rose the most between 2007 and 2019 in large CBSAs, with 41 percent of renters spending over 35 percent of their income on rent in these locations.

2.3 Migration Patterns

Figure 3 shows that rent increases have been largest for large CBSAs such as Houston and Dallas, not only in terms of average rents but also at the lower end of the rent distribution. Even the 10th percentile rents in these large metros face the steepest decline, indicating that households in traditionally lower-rent neighborhoods face substantial price pressures and that would need to move across MSAs to find lower housing costs. Yet, Figure A2 shows that the fraction of stayers is highest in large CBSAs (80%), followed by medium-sized CBSAs (70%), and lowest in small CBSAs and rural areas (50%). Among those who stayed within large CBSAs, lower-income renters are particularly immobile despite their disproportionately high rent burdens. In Houston, for example, 68% of renters with below-median-income in

2010 stayed in the same residence by 2020, 23% moved within the same neighborhood, and 9% moved to other neighborhoods within Houston. These numbers are 64%, 28% and 7%, respectively, for Dallas.

The extent to which rent increase burden incumbent renters depends not only on rent changes in their current neighborhoods, rents in alternative neighborhoods, but also on the costs of moving and the overall value of moving relative to staying. The limited mobility we observe from large CBSAs may reflect both higher cost of migration – which may be particularly burdensome for lower-income residents in large CBSAs – and the higher value of staying. For many households, remaining in a large metro like Dallas or Houston may offer economic and social benefits that outweigh the gains from relocating to areas with lower housing costs, and this value may be particularly significant for lower-income households, who often have stronger social ties, community support, or location-specific employment that anchors them in place.

Unpacking these dynamics requires understanding the value of staying versus the value of moving for different demographic groups, accountings for migration costs, the current benefits of staying versus moving, as well as the option value of moving in the future. The would help explain why, despite facing the steepest rent burdens, lower-income renters in large CBSAs remain disproportionately immobile.. Our empirical analysis is designed to estimate parameters related to these nuanced trade-offs, which can be used to evaluate how having lower moving costs and the availability of substitute neighborhoods can help mitigate the welfare costs of rapid rent growth for different income groups.

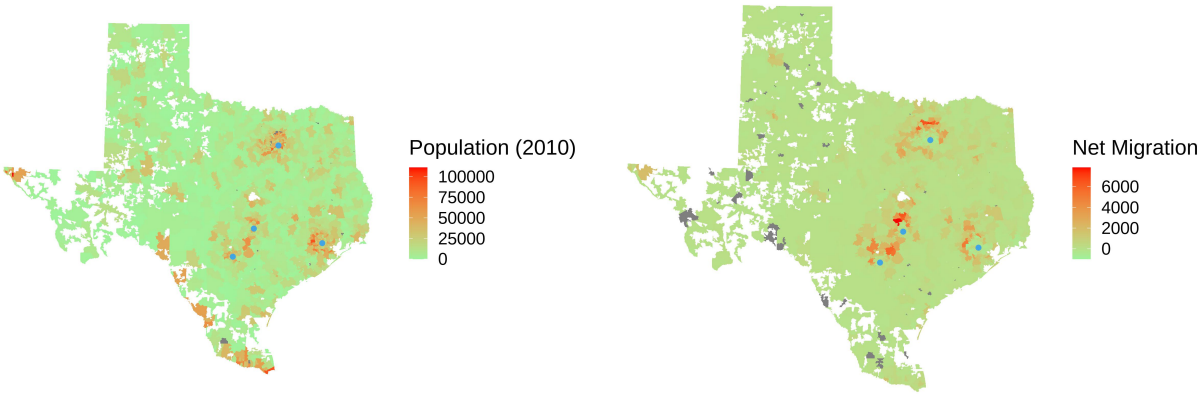
Figure 4 shows the spatial distribution of population and net in-migration in Texas, highlighting a strong concentration of movement toward Texas' four largest CBSAs: Dallas, Houston, Austin, and San Antonio. Notably, much of the population growth is occurring in urban peripheries and suburban areas, where the supply of rental units is likely more elastic and rent growth may be more moderate in response to demand shocks.

This suggests that urban fringe can become the primary margin for housing affordability policies. If new housing supply is added in these peripheral neighborhoods, where land is more available and regulatory barriers are often lower, it can help accommodate both incoming population and households relocating within the CBSA from higher-rent neighborhoods. Even if lower-income renters in high-demand neighborhoods face steep moving costs or strong community ties that discourage them from relocating to these new developments themselves, the additional supply can still mitigate rent pressure in their original neighborhoods. By absorbing demand that would otherwise spill into those neighborhoods, either from new residents moving into the CBSA or from current residents seeking lower cost alternatives, new supply in the periphery can ease overall rental market tightness.

In our empirical analysis, we will evaluate one such policy counterfactual. By simulating

the effects of adding housing supply in alternative neighborhoods, such as urban peripheries, we can examine how these changes might alleviate rent burden for households from different demographic groups.

Figure 4: Population and Migration Distribution



(a) Texas Population Map

(b) Texas Net Migration Map

Note: Net in-migration is greatest in large (Dallas, Houston) and medium size (Austin, San Antonio) CBSAs shown with blue dots on the maps.

3 Elasticities of Migration With Respect to Rents

The full dynamic model laid out in the following section features a series of nested choices. First, individuals choose whether to stay or migrate, while anticipating the expected values of moving or staying. For agents that decide to migrate, they must decide whether to migrate to another house in the same neighborhood, a different neighborhood in the same CBSA, or to a different CBSA. While the destination choice conditional on migration has a continuation value, this value is likely to be correlated with the flow utility and thus be capitalized into current rents. That is, choices of migration destinations resemble static decisions. In contrast, the decision of whether to migrate depends crucially on the continuation value in one's current location. As such, the binary out-migration choice is more of a dynamic decision.

Exploring the dynamic model's simpler static counterpart has three benefits. First, it allows us to focus on identification less opaquely than in the dynamic model. The crucial (endogenous) predictor variable of interest for us is neighborhood rent, with the key iden-

tification challenge to isolate variation in neighborhood rents that are uncorrelated with unobserved demand factors or shocks. Second, it establishes some baseline results about the relationships between rents and migration decisions, confirming the upward bias in estimates that are not corrected for unobserved heterogeneity. Finally, it demonstrates the importance of considering continuation values in the move-stay decision.

In a first exercise, we demonstrate that housing costs in potential destination neighborhoods influence location choices conditional on migrating. In a second exercise, we show that origin neighborhood rents influence the decision to migrate. Coefficients on rent in both exercises are biased upwards because of unobserved heterogeneity across neighborhoods. Neighborhoods that are more attractive tend to have higher rents, and an IV strategy is needed to correct this endogeneity problem. However, our IV strategy cannot handle correlations between continuation values and current local amenities for stayers. A strong feature of the data is reduced migration rates out of neighborhoods with rising rents, likely because these rising rents reflect these neighborhoods' growing attractiveness. Below we demonstrate the need for model structure to explicitly control for continuation values, handling this issue.

3.1 A Simple Static Model of Migration

Estimates reported in this section can be viewed in the context of the following standard static model of migration, which is a module in the full dynamic model developed in the following section.

Individual i has the indirect utility

$$\tilde{u}_{it}(l,n) = \alpha \ln R_{lnt} + \omega_{lt} + \phi_{ln} + \nu_{lnt} + \epsilon^i_{lnt} \equiv v_t(l,n) + \epsilon^i_{lnt}$$

associated with living in neighborhood n of commuting zone l in year t. With this formulation, an implicit assumption is that households can flexibly adjust the quantity of housing services they rent in each period. Following Davis et al. (2021), rent enters in logs to provide more flexibility in estimation and to be consistent with Cobb-Douglas or quasilinear utility. Utility is also influenced by commuting zone-year and neighborhood fixed effects, which reflect time varying commuting zone and static local amenities, respectively. For completeness, we also include the common unobserved neighborhood-year specific ν_{lnt}, which can be thought of as a neighborhood amenity shock. The individual level i.i.d. shock ϵ^i_{lnt} is distributed extreme value type I with dispersion parameter σ. The CBSA-year fixed effects ω_{lt} allow these neighborhood level shocks to be nested within CBSAs.[6]

[6]In one specification of the dynamic model, we allow σ to depend on CBSA attributes and observed individual characteristics.

Migrating from neighborhood n to neighborhood n' has the utility cost

$$\tau_{nn'} = \tau_{0n} + \tau_{dn} \ln d_{nn'} \mathbf{1}(l' = l) + \tau_{dl} \ln d_{nn'} \mathbf{1}(l' \neq l),$$

where $\ln d_{nn'}$ is the log distance between the two neighborhoods. We allow the moving cost function to differ flexibly for moves within versus between CBSAs. In some specifications, we also allow it to depend on observed heterogeneity. The moving cost is 0 for non-movers and it is the centripetal force that tends to keep people from moving longer distances.

Using the properties of the extreme value distribution, the number of movements to neighborhood n' conditional on migrating from origin neighborhood n is:

$$N_{n \to n'}(t) = s_{nt} \frac{\exp\left(\nu_t(l', n') - \tau_{nn'}\right)/\sigma}{\sum_{l''} \sum_{n'' \in l''} \exp\left(\nu_t(l'', n'') - \tau_{nn''}\right)/\sigma} \tag{1}$$

where s_{nt} is the population of the origin neighborhood. This implies the following standard Poisson regression form:

$$N_{n \to n'}(t) = \exp\left(\frac{\alpha}{\sigma} \ln(R_{kt}) + \frac{1}{\sigma}\tau_{nn'} + \frac{1}{\sigma}\omega_{l't} + \frac{1}{\sigma}\phi_{l'n'} + \lambda_{lnt}\right)\mu_{nn't}, \tag{2}$$

where $\mu_{nn't}$ is the error term. Origin neighborhood–time fixed effects λ_{lnt} control for origin neighborhood population and the denominator in equation (2), meaning that identifying comparisons are between rents in different destinations. Destination CBSA-time fixed effects $\omega_{l't}$ control for the general attractiveness of each CBSA in each year.[7] Destination neighborhood fixed effects $\phi_{l'n'}$ are neighborhood static amenity contributions to utility. This set of fixed effects identifies variation in $\ln R_{l'n't}$ to be within CBSA and over time. Note, that the utility parameters are normalized by the dispersion parameter σ and we do not separately identify the rent coefficient in the static model.

Many origin-destination neighborhood pairs exhibit few or zero movements in our data. One common approach to address this issue is to kernel smooth the estimated probabilities of movements across neighborhoods (Diamond et al. (2019), Schubert (2021), Almagro and Domınguez-Iino (2025), Bayer et al. (2016)). Smoothing tends to make neighborhoods look more similar, and can thus lead to biases, especially when dealing with a large number of neighborhoods and low probabilities of movement, which is the case in this study. To overcome the zero movement problem, we include all zeros and employ Poisson pseudo-maximum likelihood (PPML) to estimate parameter clusters in (2) within and across metropolitan areas (Guimaraes et al., 2003; Dingel and Tintelnot, 2020).

[7] Structural interpretations are λ_{lnt} equals log origin neighborhood population $s_{lnt} - \ln\left(\sum_{l'} \sum_{n'' \in l'} e^{(v_{it}(l', n'') - \tau_{nn''})/\sigma}\right)$, $\omega_{l't}$ controls for the time varying commuting zone amenities, and $\phi_{l'n'}$ is static neighborhood amenities.

3.2 Instrumenting for Rents and First Stage Results

We utilize an instrumental variable to identify causal effects of rental prices on migration flows. The need for an instrument arises due to the presence of unobservables that may be correlated with rents and influence neighborhood demand. Neighborhood amenities ν_{nlt} contribute to the utility of residing in neighborhood nl and are likely to be capitalized into rental prices. Hence, we require an instrument that is correlated with rental prices in a given ZCTA only through its influence on housing supply. That is, the instrument must isolate variation in $\ln R_{lnt}$ that is uncorrelated with ν_{lnt} conditional on neighborhood fixed effects and CBSA-year fixed effects.

We use a shift-share instrument based on the interaction between neighborhood housing supply elasticities and the overall nonparametric trend in Texas rental prices, which is largely driven by secular demand growth. The exposure component of the IV is the local housing supply elasticity, as estimated in Baum-Snow and Han (2024). The time-varying component of the IV is the logarithm of average rental prices in Texas. The average rental prices are calculated as the weighted average of rents across all Texas ZCTAs, with weights based on the number of housing units in each ZCTA. That is, the instrument is calculated as $Z_{lnt} = \ln \overline{R}_t \times \epsilon_{ln}$.

In response to a positive neighborhood demand shock for rental units, high-elasticity neighborhoods build more housing, resulting in a smaller increase in rental prices compared to low-elasticity neighborhoods. The identifying assumption is thus that the instrument affects rental prices only through differences in neighborhood supply-side responses to the same demand shock. As all specifications throughout our analysis control for neighborhood and CBSA-time fixed effects, threats to identification would come from differential neighborhood level trends in amenities as a function of the local housing supply elasticity. This could happen, for example, if suburban neighborhoods are on differential amenity trends from urban neighborhoods. Below we explore robustness to various such potential confounders and show that observed time-varying neighborhood amenities are not related to the instrument, conditional on controls. We note that estimated supply elasticities are constructed using census tract characteristics that are either fixed or measured as of 2000 or 2006, well prior to our study period. These characteristics are tract 2000 fraction of land area developed, plausibly exogenous components of 2000-2006 growth in residents' commuting market access beyond 2 km, measures of topography, and CBD distance. Metro area fraction of land unavailable for development in 2000 and the 2006 Wharton Land Use Regulation Index also contribute to these estimated elasticities.

Mechanically, we follow the control function approach in Lin and Wooldridge (2019) and estimate the PPML regressions with an instrument for rent. This approach proceeds in two steps. In the first step we regress the endogenous variable (log rent) on the instrument and

Table 1: First Stage Results

	Dependent Variable:			
	Log Rent			
	(1)	(2)	(3)	(4)
IV	-0.36^{***}	-0.35^{***}	-0.36^{***}	-0.36^{***}
	(0.07)	(0.07)	(0.08)	(0.08)
Controls	–	ZCTA Controls	Dist. to CBD - Year Dummy	Dist. to CBD - Year Dummy
Orig. ZCTA - Year FE	Yes	Yes	Yes	Yes
Dest. CBSA - Year FE	Yes	Yes	Yes	Yes
Dest. ZCTA FE	Yes	Yes	Yes	Yes
Clustered	Yes	Yes	Yes	Yes
IV F.Stat	27	25	22	20
Subset	Movers	Movers	Movers	All
Observations	7,538,672	7,538,672	7,353,092	7,360,224

Note: Each column is a different specification of the first stage linear regression of the log of rent on the instrument variable and indicated fixed effects. The standard errors are clustered at the origin neighborhood level. ZCTA controls in the second column are time-varying ZCTA income, population and the number of rental units, calculated using five year ACS data. Columns (3) and (4) additionally control for CBD distance by year fixed effects. Column (1)-(3) only includes movers and Column (4) additionally includes stayers.

controls to obtain predicted values and recover prediction errors. The second step performs the PPML regression and includes the prediction error from the first step as a control to account for endogeneity. Consistent standard errors in the second step are recovered through a bootstrapping procedure.

Table 1 presents first-stage regressions of rental prices on the instrument and our standard set of controls. To maintain consistency with the second-stage regression, which analyzes the movement matrix between origin ZCTA groups and destination ZCTA groups, these regressions maintain one observation per origin-destination-year triplet. Standard errors are clustered at the destination-year level.

We note that the first stage coefficient of -0.36 is very stable across specifications and samples. This coefficient has a difference in difference type interpretation. If Texas average rents increase by 10% in a year, a neighborhood with a supply elasticity of 0.75 experiences a 1.7% smaller increase in rent than does a neighborhood with a supply elasticity of 0.25. Including time-varying neighborhood demographic controls and CBD distance by year controls does not affect results. This is evidence that the instrument is not correlated with differential trends by CBD distance or with neighborhood demographics, which could generate endogenous amenities. Regressions of log ZORI or log median rents on the instrument and neighborhood and CBSA by year fixed effects run at the neighborhood by year level instead yield statistically significant coefficients of -0.32 and -0.30, respectively. Analogous regressions of neighborhood demographics on the instrument yield small insignificant coefficients (Table A4).

Table 2: Static Model PPML Regressions

Dependent Variable:	Number of Migrants			
	PPML-OLS		PPML-IV	
Model:	(1)	(2)	(3)	(4)
Variables				
Log Rent	-0.44***	-0.40***	-3.3***	-3.4***
	(0.07)	(0.05)	(0.51)	(0.54)
Log Distance		-0.51***	-0.51***	
		(0.002)	(0.002)	
Stage 1 Residual			2.9***	3.0***
			(0.51)	(0.55)
Log Distance (Same CBSA)				-0.39***
				(0.004)
Log Distance (Other CBSA)				-0.61***
				(0.003)
Fixed-effects				
Destination ZCTA Group	Yes	Yes	Yes	Yes
Origin ZCTA Group - Time	Yes	Yes	Yes	Yes
Destination CBSA - Time	Yes	Yes	Yes	Yes
Fit statistics				
Observations	7,538,672	7,538,672	7,538,672	7,531,360
Pseudo R^2	0.26632	0.69153	0.69154	0.75302

Note: The dependent variable is the number of migrants from the origin ZCTA to the destination ZCTA in each year. Columns (1) and (2) present partial correlations and columns (3) and (4) show PPML-IV results. Stage 1 Residual is the error of prediction from the first stage. "Distance" measures the distance between the origin and destination, and "Distance (Same CBSA)" is the distance interacted with a dummy variable equal to 1 if the move occurs within the same CBSA.

3.3 Second Stage Results

Table 2 shows the baseline results of the PPML regressions, which are equivalent to estimating equation (2). Regressions include controls for origin neighborhood-year, destination CBSA-year, and destination neighborhood fixed effects. Origin neighborhood-year fixed effects control for the inclusive value of migrating out of the neighborhood, shown in the denominator of equation (1). Destination CBSA-time fixed effects control for variation in the attractiveness of a CBSA, such as variation in local labor market conditions. Destination neighborhood fixed effects control for a neighborhood's static amenities and attractiveness.

The rent coefficients in columns (1) and (2) are biased toward zero because variation in rent is correlated with unobserved variation in local amenities. Additionally, reverse causality between migration and rent further biases the estimated coefficients. Columns (3) and (4) address these biases using the control function approach and PPML-IV regression. Including the prediction error from the first-stage IV regression in the second stage generates a rent coefficient of -3.3. Looking at the structural interpretation of the rent coefficient in equation (2), we see that the disutility from rent (α) cannot be separately identified from the dispersion parameter (σ). Thus, the rent coefficient represents the normalized disutility

of rent divided by the dispersion parameter. To provide context, if the rent coefficient is -0.40[8], the dispersion parameter would be approximately $0.4/3.3 \simeq 0.12$.

Column (3) in table 2 shows that a 10% increase in rent reduces the migration inflow by 33%. Moreover, the table shows that the distance between neighborhoods affects the moving cost. A 10% increase in distance reduces the migration inflow by 5.1%. However, column (4) shows that the effect of distance varies within and across CBSA migrations, with higher marginal distance costs of cross-CBSA movements.

These simple variable moving cost estimates may mask substantial heterogeneity across households. To begin unpacking this heterogeneity, Table 3 repeats the PPML-IV regressions in column 4 of Table 2 with the addition of a set of interactions between, pairwise distance, and average income, share of young households, and CBSA population of the origin neighborhood. Column 4 of Table 3 shows that, with these additional controls, the migration responds to the destination neighborhood rent and distance to the same extent as before. However, it suggests households moving from neighborhoods with a higher share of low-income renters or older populations have greater moving costs. In addition, the interaction of distance with destination CBSA population suggests both moving within and to larger CBSAs are more costly.

Tables A4 and A5 present similar results as above, but they include both movers and stayers in the regressions. For stayers, the origin and destination neighborhoods are the same, and, as expected, their elasticity with respect to destination rental prices is smaller than that of movers. Consequently, the regression that includes stayers yields a smaller rent coefficient.

3.4 Understanding Migration Decisions

Incorporated in the treatment above is the joint choice of whether to migrate and the destination conditional on migrating. In practice, this may not be a joint decision. Individuals may trade off the value of staying in their current neighborhood v_{lnt}^{stay} against the inclusive value of moving v_{lnt}^{move}, an aggregation across utilities associated with moving to all potential destinations net of moving costs. Indeed, results of regressing the neighborhood level stay probability on instrumented rents and the same fixed effects and controls as in Table 2 yields positive coefficients on rents (Table A6). These positive coefficients reflect the correlation between the instrument and continuation values for stayers.

The fully dynamic model in the following section explicitly incorporates the move-stay decision, allowing us to account for the endogeneity of the continuation value for stayers (and for movers). In particular, the move decision will depend on origin location rent and

[8]The average share of rent from total income.

Table 3: PPML-IV Regression: Heterogeneity in Moving Costs

Dependent Variable:	Number of Migrants			
Model:	(1)	(2)	(3)	(4)
Variables				
Log Rent	-3.6***	-3.6***	-3.3***	-3.5***
	(0.53)	(0.54)	(0.50)	(0.49)
Stage 1 Residual	3.2***	3.2***	3.0***	3.1***
	(0.54)	(0.54)	(0.51)	(0.50)
Log Distance (Same CBSA)	-0.38***	-0.40***	-0.34***	-0.35***
	(0.004)	(0.004)	(0.002)	(0.002)
Log Distance (Other CBSA)	-0.61***	-0.61***	-0.60***	-0.61***
	(0.003)	(0.003)	(0.002)	(0.002)
Log Distance * Low Income Share (Same CBSA)	0.13***			-0.11***
	(0.02)			(0.02)
Log Distance * Low Income Share (Other CBSA)	-0.01			-0.16***
	(0.02)			(0.01)
Log Distance * Young Share (Same CBSA)		0.31***		0.12***
		(0.04)		(0.02)
Log Distance * Young Share (Other CBSA)		0.23***		0.16***
		(0.03)		(0.02)
Log Distance * CBSA Population (Same CBSA)			-0.07***	-0.08***
			(0.001)	(0.002)
Log Distance * CBSA Population (Other CBSA)			-0.01***	-0.03***
			(0.002)	(0.002)
Fixed-effects				
Destination ZCTA Group	Yes	Yes	Yes	Yes
Origin ZCTA Group - Time	Yes	Yes	Yes	Yes
Destination CBSA - Time	Yes	Yes	Yes	Yes
Fit statistics				
Observations	7,524,048	7,524,048	7,345,960	7,338,828
Pseudo R^2	0.72176	0.72268	0.76898	0.74029

Note: The dependent variable is the number of migrants of each tenure level moving from the origin to destination neighborhood in each year. "Distance" measures the distance between the origin and destination. "(Same/Other CBSA)" indicates that the variable is interacted with a dummy equal to 1 if the move occurs within the same CBSA, and 0 otherwise. "Low Income Share" represents the share of renters with annual incomes below $35,000. "Young Share" refers to the share of renters aged 15 to 35 years. "Dest. CBSA Population" denotes the population of the destination CBSA.

fixed effects as balanced against the inclusive value of moving, which can be written as:

$$v_{lnt}^{move} = \sum_{n'} \exp\left(v_t(l', n') - \tau_{nn'}\right) \tag{3}$$

This suggests two types of estimates. First is an individual level logistic regression of moving on origin rents, CBSA-year FE, origin neighborhood FE, and v_{lnt}^{move}. Second is a regression of the $\ln \frac{movers_{lnt}}{stayers_{lnt}}$ on these same variables. In both cases, origin rents would be instrumented as above. The constant can then be interpreted as the fixed moving cost, and the coefficient on the inclusive value of moving is the moving elasticity. In the following section, we develop these ideas further using the full dynamic model.

4 Dynamic Model

Evidence in the descriptive section above indicates that migration rates are much higher within large CBSAs despite the existence of lower cost housing further afield. While movements to lower rent neighborhoods are systematic in the data, gravity frictions are quite strong, especially between CBSAs. The dynamic model is oriented toward allowing us to distinguish between the following mechanisms that may explain these observations. First, migration costs may deter moving to better opportunity locations. That is, we need estimates of the migration technology that are flexible enough to handle both within and cross CBSA moves by individual type. Second, the availability of alternative neighborhoods that are close substitutes or preferred to current neighborhoods in the same or other CBSAs may be limited. That is, we need estimates of the demand system for neighborhoods within and between CBSAs. Third, there may be fixed moving costs that are random. These would capture idiosyncracies about landlord-tenant relationships, shocks to housing unit quality, and shocks to family conditions including marriage, divorce and having children. Fourth, individuals may get value out of ties to their current home and CBSA, which may develop over time. The structure of the dynamic model allows us to measure the magnitudes of these third and fourth forces. Moreover, the dynamic model helps with identification. It allows us to control for the possibility that agents value the future expected option value associated with neighborhoods beyond the period flow utility. As discussed above, this option value may be correlated with our instrument for rents.

We approach these aims by specifying a nested model under two alternative specifications. The top nest, common to both specifications, is the move-stay decision. Here, individuals move if the composite value of moving to an alternative home net of migration costs, with one option in each neighborhood, exceeds the flow utility associated with staying in the home plus a continuation value. Observed variation in rates of neighborhood choices (conditional choice probabilities) made by migrants from each origin in each year provide sufficient information to allow us to estimate the conditional values of each destination net of migration cost, through the standard Hotz and Miller (1993) inversion. As the resulting choice-specific values from each origin incorporate both the value of living in the new neighborhood and the migration cost between neighborhoods, we decompose the two components using gravity relationships. In a first model specification, we estimate all values conditional on moving jointly with the move-stay decision using a neural net. As in Bodéré (2023), value functions are flexibly parameterized as a function of individual, origin neighborhood, and destination neighborhood observables. In the second specification, we separately estimate the conditional renewal values of migrating to each neighborhood and migration costs as an aggregated nested logit PPML high dimensional fixed effects gravity regression. While this second specification requires partitioning into fewer discrete types, it is arguably more

transparent, with superior robustness properties.

In both specifications, moving costs are specified to flexibly differ within and between commuting zones as a function of individual characteristics, tenure, and migration distance. Elasticities of substitution in demand across neighborhoods are specified to flexibly differ within and between different commuting zones of different sizes. We also accommodate local neighborhood amenities and marginal (dis)utilities of rent. As all parameters and values are potentially indexed by individual level observables Z_{it}, we suppress this explicit dependence in the notation for the remainder of this section.

4.1 Preferences and Migration Costs

Expanding on the specification in the prior section, the deterministic component of the flow utility an agent enjoys from living in commuting zone l neighborhood n in year t is

$$u_t(l, n, \iota^l, \iota^r) = \alpha \ln R_{lnt} + \beta_{\iota l}\iota^l + \beta_{\iota r}\iota^r + \omega_{lt} + \phi_{ln} + \nu_{lnt}. \tag{4}$$

Flow utility depends on the neighborhood rent index R_{lnt}, tenure in the commuting zone ι^l, tenure in the home ι^r, a commuting zone-year fixed effect ω_{lt}, a neighborhood fixed effect ϕ_{ln}, and a time varying neighborhood component ν_{lnt}. These final three components account for local amenities and labor market opportunities and may be correlated with rents. Relative to the static model formulation, we additionally allow tenure in the home and CBSA to contribute to flow utility in order to make sure that the estimated model can separate out tenure effects on preferences from moving costs. We want to allow individuals to gain stronger ties to their homes and CBSAs over time.[9]

Each year, households can either stay in their apartment or leave their apartment for another home. As the estimation sample only includes those who remain renters throughout the period we observe them, we avoid the need to incorporate portfolio choice and asset accumulation in the model. The moving cost to a new neighborhood in the same commuting zone, $\tau^s_{nn'}(\iota^r)$, depends on the distance between neighborhood n and n' and the time spent living in the home prior to the move, ι^r. The moving cost to a new neighborhood in a different commuting zone is $\tau^o_{nn'}(\iota^l, \iota^r)$, and additionally depends on the time spent living in the origin (l) commuting zone, ι^l. The fixed moving cost, paid for all types of moves, is $\tau^s_0(\iota^r)$. These moving costs are paid immediately after moving. Those that stay in their

[9] In a similar setup, Howard and Shao (2023) incorporates spatial autocorrelation in amenities to account for ties to origin locations. See Pakes et al. (2021) for an alternative approach to dealing with this issue in the context of health plan choice.

apartment pay no moving cost. Putting this together, we specify moving costs as:

$$\tau_{nn'}(\iota^l, \iota^r) = \begin{cases} \tau^s_{nn'} = \tau^s_0(\iota^r) + \tau^s_d(\iota^r) \ln d_{nn'} & \text{if } n' \in l \\ \tau^o_{nn'} = \tau^s_0(\iota^r) + \tau^o_0(\iota^l) + \tau^o_d(\iota^l) \ln d_{nn'} & \text{if } n' \in l' \neq l \end{cases} \tag{5}$$

In the empirical implementation, we make the indicated functions $\tau_0(\cdot)$ and $\tau_d(\cdot)$ linear in their arguments.

4.2 Value Functions

The full value function associated with being in neighborhood ln with CBSA and home tenures ι^l and ι^r, respectively, at the beginning of period t is the maximum over two options. The first option is to stay in the same home for one additional year, which comes with the flow utility $u_t(l, n, \iota^l, \iota^r)$ plus an individual-specific idiosyncratic shock ϵ^s_{ti} plus the discounted expected value of starting period $t+1$ in neighborhood ln. The second option is to migrate to an alternative home, for which there is one idiosyncratic draw in each neighborhood, minus any moving costs. At the beginning of each period, agents see each move-neighborhood-location-specific shock ϵ^m_{lnti} and decide whether and where to move by choosing the highest conditional value option. In the neural net specification, agents have a standard logit demand system across these options, with variable substitution elasticities across choices driven by variation in observable individual and neighborhood characteristics. In our nested logit specification, this shock is nested for the commuting zone and neighborhood within the commuting zone conditional on moving, as in Couture and Handbury (2020).

The upper "move decision" nest imposes a correlation between all shocks associated with living in a different home. There are a total of $N = \sum_{l=1}^{L} N_l$ neighborhoods in the choice set, with each commuting zone l having N_l neighborhoods and N_l shocks. One additional shock is associated with remaining in the same home. This home-specific shock represents either landlord choices or changes in family or employment that idiosyncratically change the match quality between homes and individuals. The control variables are commuting zone l' and neighborhood n' next period, which can remain the same if the agent does not move or moves but stays in the same neighborhood. The state variables are commuting zone last period l, neighborhood last period n, time lived in the current commuting zone and residence, ι^l and ι^r respectively, exogenous observables Z_{it}, and flow utility shocks.

The value in year t of being in neighborhood ln with ι^l years in the CBSA and ι^r years in the home is

$$V_t(l, n, \iota^l, \iota^r) = \max\{V^{stay}_t(l, n, \iota^l, \iota^r), V^{move}_t(l, n, \iota^l, \iota^r)\} \tag{6}$$

The value of staying is

$$V_t^{stay}(l, n, \iota^l, \iota^r) = u_t(l, n, \iota^l, \iota^r) + \epsilon_{ti}^s + \beta \bar{V}_{t+1}(l, n, \iota^{l\prime}, \iota^{r\prime}). \tag{7}$$

The value of moving is

$$\begin{aligned} V_t^{move}(l, n, \iota^l, \iota^r) = \max\{ & \\ & \max_{n'|l} [u_t(l, n', \iota^l, 0) - \tau_{nn'}^s(\iota^l, \iota^r) + \epsilon_{ln'ti}^m + \beta \bar{V}_{t+1}(l, n', \iota^{l\prime}, 1)], \\ & \max_{n'|l' \neq l, l'} [u_t(l', n', 0, 0) - \tau_{nn'}^o(\iota^l, \iota^r) + \epsilon_{l'n'ti}^m + \beta \bar{V}_{t+1}(l, n', 1, 1)]\}. \end{aligned} \tag{8}$$

This is the maximum over the value of staying in each possible migration destination, with ι^r reset to 0 for all moves and ι^l reset to 0 for moves across CBSAs, net of migration costs and for alternative draws of idiosyncratic shocks. That is, the resetting of tenure lengths to 0 makes the value of moving depend on the "renewal values" $V_t^{stay}(l, n, \iota^l, 0) - \epsilon_{ti}^s + \epsilon_{ln'ti}^m$ if moving within CBSA l and $V_t^{stay}(l, n, 0, 0) - \epsilon_{ti}^s + \epsilon_{l'n'ti}^m$ if moving between CBSAs (Arcidiacono and Miller, 2011). This is a convenient feature that simplifies estimation.

The ex-ante value function $\bar{V}_{t+1}(\cdot) = \mathbb{E}_t V_{t+1}(l, n, \iota^l, \iota^r)$ is

$$\bar{V}_{t+1}(l, n, \iota^l, \iota^r) = \ln\{\exp \bar{V}_{t+1}^{stay}(l, n, \iota^{l'}, \iota^{r'}) + \exp \bar{V}_{t+1}^{move}(l, n, \iota^{l'}, \iota^{l'})\} + \Gamma. \tag{9}$$

It takes this form given the extreme value type I shock distribution. The ex-ante value of staying is $\bar{V}_t^{stay}(l, n, \iota^l, \iota^r) = V_t^{stay}(l, n, \iota^l, \iota^r) - \epsilon_{ti}^s$. The corresponding ex-ante value of moving is $\bar{V}_t^{move}(l, n, \iota^l, \iota^r) = -\tau_0^s(\iota^r) + \gamma IV_t^m(l, n, \iota^l, \iota^r)$, where the inclusive value of moving $IV_t^m(l, n, \iota^l, \iota^r)$ is specified below, with its specific form depending on the joint distribution of moving shocks ϵ_{lnti}^m across neighborhoods ln for each individual i in each year t. The fixed migration cost $\tau_0^s(\iota^r)$, also seen in equation (5), factors out of this term, and so is a remaining constant. Γ is Euler's constant.

The inclusive value of moving, $IV_t^m(l, n, \iota^l, \iota^r)$, summarizes the attractiveness of all moving options. It can be thought of as a summary measure of the attractiveness of alternative neighborhoods given the state vector and is a measure of migration market access. It will thus tend to be higher in location-years in which higher value neighborhoods are available at lower migration costs. In the neural net specification, it takes the following form:

$$IV_t^m(l, n, \iota^l, \iota^r) = \ln\left(\sum_{l'=1}^{L} \sum_{n' \in N_{l'}} \exp\left[V_t^{nn'}(\iota^l, \iota^r, Z_{it})\right]\right) \tag{10}$$

In this expression, $V_t^{nn'}(\iota^l, \iota^r, Z_{it})$ denotes the value of being a new arrival in location n',

having migrated from n. It is a function of observables about the individual and the neighborhoods, with the exact set of predictive features employed specified in the following section. The structural interpretation of $V_t^{nn'}(\iota^l, \iota^r, Z_{it})$ is

$$V_t^{nn'}(\iota^l, \iota^r, Z_{it}) = [\bar{V}_t^{stay}(l', n', \iota^{l'}, 0, Z_{it}) - \tau_{nn'}(\iota^l, \iota^r, Z_{it}) + \tau_0^s(\iota^r, Z_{it})]/\gamma. \tag{11}$$

That is, the pairwise migration value can be decomposed into a renewal value and a variable migration cost, both of which are scaled by the Extreme Value Type I distribution dispersion parameter γ. The fixed migration cost $\tau_0^s(\iota^r, Z_{it})$ factors out of the inclusive value of moving, as it is identical for all origin-destination pairs. Here, $\iota^{l'} = \iota^l$ for migrants within CBSAs and $\iota^{l'} = 0$ for migrants between CBSAs. Equation (20) in Appendix A specifies the corresponding inclusive value of moving in the nested logit model specification.

A central parameter governing the move-stay decision is the logit dispersion parameter γ. As γ rises, the correlation between neighborhood choice-specific values declines, thereby making moving more likely to be attractive through a high error draw. Therefore, γ can be interpreted as a general "preference" for moving. Alternatively and equivalently, γ can be viewed as a push factor out of one's home. This could be the rate of exogenous separation due to landlords' lease non-renewals or the rate at which family structure changes, prompting idiosyncratically low values of staying.

5 Estimation Methodology

Estimation of this model follows ideas in Arcidiacono and Miller (2011), Scott (2014), and Diamond et al. (2018), as described in detail in Kalouptsidi et al. (2021), and Bodéré (2023). However, because of the high dimensional state space we estimate the model in four stages in a process with some similarities to that in French et al. (2023). Using data on migrants only, we first generate predicted individual-specific conditional neighborhood choice probabilities $p_t^{nn'}(\iota^l, \iota^r, Z_{it})$ using a neural net machine learning prediction model, with corresponding migration values $V_t^{nn'}(\iota^l, \iota^r, Z_{it})$ in equation (11) easily recovered given the multinomial logit specification of the neural net. In a second step, we decompose these composite values into migration cost and renewal value components using a gravity regression. In a third stage of estimation, the move-stay decision is used to recover components of flow utilities, with the inclusive value of moving calculated in the first step as a key regressor. IV regressions on an Euler equation recover marginal utility parameters and the dispersion parameter γ. In a final step, fixed moving costs are estimated by looking at move-stay decisions among those who move within neighborhoods.

We note that a commonly used alternative approach for the first two stages of estimation

is to invert nonparametrically defined conditional choice probabilities to back out conditional values (e.g. Bayer et al., 2016 and Diamond et al., 2018). In our setting, the use of the neural net has the advantage of obviating the requirement to explicitly smooth probabilities across the type and state space in the first step. Instead, the neural net makes efficient use of observables to achieve the same goal. Our alternative nested logit formulation, explained in Appendix A, achieves this goal through parametric assumptions and a discretized state space, allowing for aggregation of migration flows and the use of a gravity regression to recover conditional values.

5.1 Specification of the Neural Net

We set up a neural net to estimate the migration value function $V_t^{nn'}(\iota^l, \iota^r, Z_{it})$ for each origin state-destination neighborhood pair and each individual migrant in the data set. This result is a parameterized value function that can be evaluated to return the migration value to each neighborhood for any state vector. We emphasize that this is purely a prediction model with no causal interpretation available. The neural net is simply a flexible way of constructing conditional values that efficiently uses observable information about individuals and their origin neighborhoods.

We use a feedforward classification neural network for this exercise. The classification model has 1030 classes, where each class represents a destination neighborhood ln. The model is defined at the migrant-year level and implemented using a two-layer architecture. Inputs to the model include a learned embedding of the origin location and a vector of pre-processed features capturing characteristics of the individual and their origin ZCTA. These inputs are concatenated and passed through fully connected layers to generate a scalar value score for each destination alternative. More specifically, the model takes as inputs: (1) the origin ZCTA indicator along with the latitude and longitude of its centroid; (2) year fixed effects for 2011-2019 and a continuous time trend; (3) 2010 origin ZCTA characteristics including median income, share of households with children, total population, share of adults with a bachelor's degree, and homeownership rate; (4) individual characteristics in the Infutor data, including household income, age at move, and an indicator for having children;[10] and (5) years in the current residence ι^r and years in the current CBSA ι^l. The model is trained to maximize the log-likelihood of observed location choices.

Because the model must accommodate 1,894 distinct origin ZCTAs, prediction using a separate dummy variable for each ZCTA would be both computationally inefficient and statistically unwieldy. Instead, we introduce an embedding layer that maps the origin ZCTA code into a 16-dimensional dense vector. This approach allows the model to learn latent

[10] As individual characteristics are only observed for about half of Infutor migrants, we impute for others using the average values from observed individuals in each neighborhood-year.

spatial characteristics that capture similarities and differences in migration patterns across nearby origins, while keeping the parameter space tractable. The resulting 16-dimensional origin embedding is concatenated with the vector of covariates laid above of length $d = 20$, producing a combined feature vector of dimension (16 + d). This input is passed through a feedforward neural network composed of two fully connected layers. The first linear layer projects the input to a 64-dimensional hidden representation, followed by a ReLU activation function and a dropout layer with dropout probability 0.1, which serves to regularize the model and reduce overfitting. The second linear layer outputs a vector of length equal to the number of destination alternatives N (ZCTA groups), producing utility scores for each option.

The architecture consists of one embedding layer and a sequential network of two linear layers, one non-linear activation, and one dropout layer. The total set of trainable parameters includes: the origin embedding matrix (of dimension number of ZCTAs $\times$ 16), the weight matrix and bias for the first linear layer ((16 + d) $\times$ 64 + 64 parameters), and the weight matrix and bias for the second linear layer (64 $\times$ J + J, where J is the number of neighborhood destinations. The output is then passed through a softmax normalization to predict the probability of selecting each destination and estimate the log-likelihood objective function.

The training data consists of a randomly selected 80% sample of the full dataset of migrants. The remaining 20% is held out as a test set and used exclusively for model evaluation. No information from the test set is used during training to ensure an unbiased assessment of out-of-sample predictive performance. The model minimizes a cross-entropy loss between the predicted destination probabilities and the true observed destination of movers. We use the AdamW optimizer (optim_ adamw), which is an extension of the widely used Adam algorithm. The optimizer updates all trainable parameters of the model, including the embedding weights and the weights and biases of both fully connected layers.

5.2 Estimation Stage 2: Migration Cost Parameters and Substitution Elasticities

5.2.1 Migration Costs

We decompose the conditional migration values estimated by the neural net into components of variable migration costs and using a gravity regression. This allows us to separate out the migration costs from the renewal values (both scaled by γ). In particular, (11) is decomposed

with the following regression:

$$\begin{aligned} V_t^{nn'}(\iota^l, \iota^r, Z_{it}) =& \mathbf{1}(l' = l)[(\tilde{\tau}_{dn}(Z_{it}) + \tilde{\tau}_{\iota r}\iota^r) \ln d_{nn'}] \\ &+ \mathbf{1}(l' \neq l)[\tilde{\tau}_{0l}(\iota^l, Z_{it}) + (\tilde{\tau}_{dl}(Z_{it}) + \tilde{\tau}_{\iota l}\iota^l) \ln d_{nn'}] \\ &+ g(\iota^r, Z_{it}) + \lambda^1_{lnt} + \lambda^2_{l'n't} + u_{lnti} \end{aligned} \tag{12}$$

This gravity regression, with origin and destination neighborhood-year fixed effects λ^1_{lnt} and $\lambda^2_{l'n't}$ is run at the individual level, with one observation per individual-year.[11] Destination fixed effects capture the renewal values enjoyed by all migrants to destination neighborhood ln in year t $\bar{V}_t^{stay}(l', n', \iota^{l'}, 0, Z_{it})$. We parameterize renewal values to be linear in ι^l, so that $\tilde{\tau}_{0l}(\iota^l, Z_{it})$ captures a combination of the fixed costs of moving across CBSAs and differences in renewal values for cross CBSA versus within CBSA migrants (who keep $\iota^l > 0$ after moving). $\tilde{\tau}_{0l}(\iota^l, Z_{it})$ can thus be thought of as capturing the full implicit cost of moving across CBSAs divided by γ. (We do not use these estimated destination fixed effects in further estimation steps.) Tildes over τ parameters denote the negative of their counterparts in equation (5) divided by γ.

Two additional observations are of note. First, distance appears in this estimating equation even though it is not a predictive feature used by the neural net. The machine learning algorithm learns that distance matters for migration. As distance is not randomly assigned, causal interpretation of distance coefficients requires that the origin and destination fixed effects control for unobserved heterogeneity and renewal values. As renewal values are likely to be similar for people coming from different origin neighborhoods, it is reasonable to treat distance coefficients as capturing structural variable migration cost parameters. Second, if all relevant individual heterogeneity were observed in characteristics Z, the origin fixed effects would not be needed. They are only included to account for potential unobserved heterogeneity.

5.2.2 Substitution Elasticities

The use of the neural net to generate conditional values allows for flexible substitution patterns across neighborhoods as a function of observables. In Appendix A, we specify and estimate a nested logit model with more model parameters as an alternative approach to the same goal. To characterize neighborhood substitution patterns delivered by the neural net, we regress predicted individual choice probabilities on predicted CBSA-time specific inclusive values, which are increasing in renewal value dispersion, while controlling for migration costs and destination CBSA fixed effects. The responsiveness of choice probabilities to inclusive

[11] While we calculate individual level migration values for all pairs of neighborhoods in each year, we only use the chosen neighborhood for this regression.

values summarizes the preference for neighborhood variety, or cross-neighborhood substitution elasticities. We characterize parameters summarizing these neighborhood substitution elasticities as a function of CBSA and individual characteristics. We also use them as a basis for comparison to parameter estimates from the nested logit model specification.

The nested logit model specified in Appendix A delivers equations (23) and (24) that relate conditional neighborhood choice probabilities to CBSA inclusive values. In our neural net model specification, these inclusive values are defined as

$$IV_{lnl't}^{NN}(\iota^l, \iota^r, Z_{it}) = \ln\left(\sum_{n' \in N_{l'}} \exp\left[Y_t(l'n', \iota^{l'}, 0, Z_{it}) - \tau_{nn'}^{v}(\iota^l, \iota^r, Z_{it}\right]/\gamma\right), \tag{13}$$

where $Y_t(l'n', \iota^{l'}, 0, Z_{it})$ is the renewal value demeaned within CBSAs and $\tau_{nn'}^{v}(\iota^l, \iota^r, Z_{it}$ is the variable migration cost. See Appendix A for a full derivation in the context of the nested logit model formulation. Inclusive values are aggregations of adjusted renewal values net of migration costs across all destination neighborhoods in each CBSA. Because of variation in migration costs to each destination by state, including origin neighborhood, and differences in renewal values by individual characteristics Z, inclusive values also differ by origin state.

Estimated parameters in the following gravity regression are used to characterize substitution patterns across migration destinations.

$$\begin{aligned}
\ln p_t^{nn'}(\iota^l, \iota^r, Z_{it}) = \\
&\mathbf{1}(l' = l)[(\tilde{\tau}_{dn}(Z_{it}) + \tilde{\tau}_{\iota r}\iota^r)\ln d_{nn'}] + \tilde{W}_{lt}(\iota^l, 0) + \tilde{\sigma}^s(X_l, Z_{it})IV_{lnlt}^{NN}(\iota^l, \iota^r, Z_{it})] \\
&+\mathbf{1}(l' \neq l)[\tilde{\tau}_{0l}(\iota^l, Z_{it}) + (\tilde{\tau}_{dl}(Z_{it}) + \tilde{\tau}_{\iota l}\iota^l)\ln d_{nn'}] + \tilde{W}_{lt}(0, 0) \\
&+\tilde{\sigma}^o(X_{l'}, Z_{it})IV_{lnl't}^{NN}(\iota^l, \iota^r, Z_{it})] + h(\iota^r, Z_{it}) + \Lambda_{lnt}^1 + \Lambda_{l'n't}^2 + v_{lnti}
\end{aligned} \tag{14}$$

This regression resembles equation (12), except that it includes additional controls for destination CBSA fixed effects and CBSA-specific inclusive values interacted with moves within or between CBSAs.

The central parameters of interest are $\tilde{\sigma}^s(X_l, Z_{it}) \equiv \sigma^s(X_l, Z_{it}) - 1$ and $\tilde{\sigma}^o(X_{l',Z_{it}}) \equiv \sigma^o(X_{l'}, Z_{it}) - 1$. $\sigma^o(X_{l'}, Z_{it})$ is the parameter governing shock dispersion in nest l' for those that migrate to l' from a different CBSA. We allow it to depend linearly on CBSA characteristics $X_{l'}$ and individual characteristics Z_{it}. $\sigma^s(X_l)$ is the corresponding parameter governing shock dispersion across neighborhoods for those migrating within CBSA l.[12] The correlation of the error term across neighborhoods within CBSA l is close to $1 - \sigma_l$. A higher correlation implies that neighborhoods are closer demand substitutes.

[12] We verify that $\sigma_l^s > 0$ and $\sigma_l^o > 0, \forall l$, which are required in order for neighborhood choice probabilities to be increasing in their values.

5.3 Estimation Stage 3: Flow Utilities and $\bar{V}^{stay}$

Now that all choice-specific value functions have been obtained, we turn to isolating components of flow utilities $u_t(l, n, \iota^l, \iota^r)$ in equation (4) and estimating the dispersion parameter γ. This amounts to estimating the top nest of the nested logit – the move-stay decision. This decision depends on the relative attractiveness of the migration option, summarized in the inclusive value of moving, and the value of staying in the same home for one more year. Conditional on the value of moving, higher stay rates indicate greater values of staying. γ summarizes the responsiveness to migration access.

The probability of staying in the same residence for year t can be written as

$$p_{lnt}^{stay}(\iota^l, \iota^r, Z_{it}) = \left(\frac{\exp\left[\bar{V}_t^{stay}(l, n, \iota^l, \iota^r, Z_{it}) + \tau_0^s(\iota^r, Z_{it})\right]}{\exp \gamma IV_t^m(l, n, \iota^l, \iota^r, Z_{it}) + \exp\left[\bar{V}_t^{stay}(l, n, \iota^l, \iota^r, Z_{it}) + \tau_0^s(\iota^r, Z_{it})\right]} \right) \quad (15)$$

This form implies that a simple logistic regression can recover $\bar{V}_t^{stay}(l, n, \iota^l, \iota^r, Z_{it}) + \tau_0^s(\iota^r, Z_{it})$ and γ given estimates of $IV_t^m(l, n, \iota^l, \iota^r, Z_{it})$ from the first step. We now have estimates of $\bar{V}_t^{stay} + \tau_0^s$, IV_t^m, and γ and can calibrate $\beta = 0.9$. Use the fact that the flow utility equals the value of staying minus the discounted continuation value. A resulting structural equation is (omitting explicit state dependencies except year for notational simplicity)

$$u_t + (1 - \beta)\tau_0^s = [\bar{V}_t^{stay} + \tau_0^s] - \beta \ln\left(\exp\left[\bar{V}_{t+1}^{stay} + \tau_0^s\right] + \exp\left[\gamma IV_{t+1}^m\right]\right) + \xi_{it} \quad (16)$$

That is, with available information, we can build the flow utility plus a rescaled fixed moving cost.

We also must face the issue that Infutor under-counts long-term stayers. To address this issue, we match the distribution of renter tenure ι^r in each ZCTA-year in the ACS data to that in the Infutor data. We assume that Infutor captures all movers correctly but misses some stayers. For this reason, we upweight stayers in the data, with weights constructed using ACS data.

We undertake two complementary approaches to estimate the coefficient on log rent in flow utility $\alpha(Z_{it})$ and the moving preference γ. The first uses a neural net to estimate $\bar{V}_t^{stay}(l, n, \iota^l, \iota^r, Z_{it}) + \tau_0^s(\iota^r, Z_{it})$ and γ using individual data and the estimated inclusive value of moving from Step 2. Then, flow utility is decomposed as the difference between $\bar{V}_t^{stay}(l, n, \iota^l, \iota^r, Z_{it})$ and its continuation value. A second approach, explained in Section A.5 uses data aggregated to the neighborhood-year-ι level based on the PPML estimates. Both approaches yield similar results.

Our primary approach uses a flexible neural net specified with the same predictive features as described in Section 5.1 and ZCTA log rents plus estimates of $IV_t^m(l, n, \iota^l, \iota^r, Z_{it})$ to classify stayers and parameterize $\bar{V}_t^{stay}(l, n, \iota^l, \iota^r, Z_{it}) + \tau_0^s(\iota^r, Z_{it})$. That is, we use a neu-

ral net classifier to predict each individual's move-stay decision in each year. This leads to individual-specific estimates of the probability of staying backed out from and a γ estimate of 0.20. Using equation (15), we can then calculate individual-specific estimates of $\bar{V}_t^{stay}(l, n, \iota^l, \iota^r, Z_{it}) + \tau_0^s(\iota^r, Z_{it})$. As we calculate predictions of $\bar{V}_t^{stay}(l, n, \iota^l, \iota^r, Z_{it}) + \tau_0^s(\iota^r, Z_{it})$, $\bar{V}_{t+1}^{stay}(l, n, \iota^l + 1, \iota^r + 1, Z_{it}) + \tau_0^s(\iota^r + 1, Z_{it})$ and $IV_{lnt+1}^m(\iota^l + 1, \iota^r + 1, Z_{it})$ for each individual in the data, all movers and stayers are used to estimate components of flow utilities.

Following the form in equation (4), we regress these flow utilities on log rent, ι^l, ι^r, neighborhood, and CBSA-year fixed effects, instrumenting for rents as in Section 3.2 above.

5.4 Estimating Fixed Moving Costs

Fixed moving costs are estimated jointly with the value of staying from the move-stay decision. To recover estimates of fixed moving costs $\tau_0^s(\iota^r)$, we use the fact that movers to neighborhood ln and stayers in neighborhood ln must have similar values of staying. Given that we have estimates of the variable components of moving costs, some differencing of values net of variable moving costs and probabilities lets us recover the fixed moving costs. In a complementary exercise, we compare only those moving within a neighborhood to stayers, as the former pays 0 variable moving cost.

In the third step of estimation, we recovered values of staying plus fixed moving costs. In the first step of estimation, we recovered renewal values. Combining these two objects for different individuals i and i' with the same characteristics Z, we have

$$[\bar{V}_t^{stay}(l', n', \iota^l, \iota^r, Z_{it}) + \tau_0^s(\iota^r, Z_{it})] - \bar{V}_t^{stay}(l', n', \iota^{l'}, 0, Z_{i't}). \tag{17}$$

Equation (17) highlights the main challenge in identifying fixed moving costs. ι^r will always be greater for neighborhood incumbents than new arrivals. Therefore, fixed moving costs are not nonparametrically separately identified from values.

Inferring fixed moving costs thus requires choosing either to interpolate what $\bar{V}_t^{stay}(l', n', \iota^l, 0, Z_{it})$ would be for people who are not new arrivals from estimates of $\bar{V}_t^{stay}(l', n', \iota^l, \iota^r, Z_{it})$ for $\iota^r > 0$ through flexible controls for ι^r (conditional on ι^l and Z) or through assuming that $\bar{V}_t^{stay}(l', n', \iota^l, 1, Z_{it}) = \bar{V}_t^{stay}(l', n', \iota^l, 0, Z_{it})$. If the value of staying is sufficiently similar for new arrivals as for incumbents with short tenures in a neighborhood, these values difference to 0 and we are left with the fixed moving cost only in equation (17). Another issue that can matter for recovering fixed moving cost estimates is the credibility of the identification of $\bar{V}_t^{stay}(l', n', \iota^l, 0, Z_{it})$. This requires residualization of variable moving costs in Step 2. One way to circumvent this issue is to only use movers within neighborhoods to identify fixed moving costs.

With these difficulties in mind, we estimate fixed moving costs in two ways. Our primary approach is to calculate the expression in Equation (17) using all moves in our data only. We restrict to movers so as to keep the sample the same across components of moving cost estimation. For this sample, we first use the neural net fit in Step 3 to predict $\bar{V}_t^{stay}(l, n, \iota^l, 0, Z_{it}) + \tau_0^s(0, Z_{it})$. The neural net's parameterization of the value function extrapolates to allow for calculation of the value of staying plus fixed moving cost given ι^r is set to 0. To assess overfitting, we also evaluate using the analogous object but predicted for each individual's actual ι^r prior to moving, which is fully in sample. Then, using the neural net from Step 1 we predict $\bar{V}_t^{stay}(l, n, \iota^{l'}, 0, Z_{it})$, which is the renewal value for moves within the neighborhood. Finally, following equation (17), we regress the difference in these objects on a constant and characteristics ι^l and Z_{it}. For the case in which we predict out for $\iota^r > 0$, we also control for ι^r.

A similar and complementary approach is to compare the unconditional probability of moving within neighborhood ln to the probability of staying in that neighborhood, as in the following equation.

$$\ln p_t^{nn'}(\iota^l, \iota_m^r, Z_{it}) + \ln \frac{p_{lnt}^{move}(\iota^l, \iota_m^r, Z_{it})}{p_{lnt}^{stay}(\iota^l, \iota_s^r, Z_{it})} = -\tau_0(Z_{it}, \iota_s^r) - \frac{\Delta \bar{V}_t^{stay}(l, n, \iota^l, \iota^r, Z_{it})}{\Delta \iota^r} \iota_s^r \tag{18}$$

In this equation, ι_m^r is home tenure for movers who are compared to stayers with tenure ι_s^r. The log ratio of moves within the same neighborhood to the number of stays delivers the fixed cost of moving plus the difference in values, which we parameterize as linear such that $\frac{\Delta \bar{V}_t^{stay}(l,n,\iota^l,\iota^r,Z_{it})}{\Delta \iota^r}$ is a constant. Intuitively, if the fixed cost of moving is small, we expect a higher number of moves within neighborhoods. We implement equation (18) as a regression. This aggregate approach has the advantage of using actual rather than predicted moves and stays for identification but the disadvantage of being able to accommodate less unobserved heterogeneity.

6 Parameter Estimates

Before reporting other parameter estimates, it is instructive to preview estimates of the marginal utility of log rent $\alpha(Z_{it})$ and the migration preference parameter γ. With these, we can convert all other relevant parameters to dollar terms. The marginal utility of log rent is estimated to average about -0.3, with more negative marginal utilities for lower income tenants, older tenants, and those with children. Given this, other marginal values can be put in willingness to pay terms in dollars by multiplying them by (annual rent)/0.3. This is about \$56,000 for the average rental unit in large and medium sized CBSAs and \$40,000 for the average rental unit in small CBSAs and rural areas (Figure 3). The resulting interpretation

after this dollar conversion is thus the one year rent discount required to make individuals indifferent to a one unit change in the underlying parameter.

The Step 3 neural net estimates γ to be 0.2. While we use this estimate in our calculations, our alternative aggregate regression method returns an estimate of 0.4. We elect to stick with the lower estimate as it comes from employing the same methods as the other parameters on which we focus. Future work will investigate this discrepancy further.

6.1 Variable Moving Costs

We begin with the decompositions of movers' pairwise values given in equation (11) using gravity regression specified in equation (12). Each of the reported variable moving cost coefficients should be multiplied by $\gamma = 0.2$ to recover underlying τ parameters in the moving cost function in equation (5). Therefore, putting them in dollar terms requires multiplying by about $\$10,000$.

Table 4 shows the results. Moving from left to right, we add various interactions and controls, with Columns (5) and (6) having the most comprehensive set of marginal migration cost estimates. Column (6) includes more finely defined origin fixed effects, in which origin neighborhood-year is interacted with above and below median Infutor reported demographics, to control more finely for individual level heterogeneity. As all predictor variables are demeaned except tenure lengths, coefficients apply to the average migrant with tenure lengths of 0. Stability if coefficients across specifications is evidence that unobserved heterogeneity is unlikely to be driving estimates.

Estimates in Table 4 provide considerable evidence that cross-CBSA moves are much costlier those within CBSAs. Coefficients in the top row show that, for average distance migrants, it is about $65,000 more costly to migrate between than within a CBSA. Doubling migration distance within a CBSA costs about $2,000 more. Between CBSAs, doubling migration distance is about $10,000 costlier. Heterogeneity in migration costs by length of stay in the CBSA are also evident. Staying longer in a CBSA increases the migration cost out of that CBSA (fourth row). It also increases the cost of migrating further (sixth row). However, the same patterns are much smaller for longer tenure in the home.

Rows with "Orig" demographics indicate heterogeneity by 2010 attributes of origin neighborhoods. Migrants from higher income neighborhoods are much less sensitive to cross-CBSA moves and distance. Migrants from neighborhoods with high child shares are less sensitive to distance for within-CBSA moves. Remaining indicated demographics use individual level information reported in Infutor, with similar patterns. Corresponding variable migration cost estimates using aggregate PPML gravity regressions derived from the nested logit model with a discretized state space are qualitatively similar, and reported in Table

Table 4: Neural Net Variable Moving Cost Estimates

	Dependent variable:					
	Destination Value					
	(1)	(2)	(3)	(4)	(5)	(6)
Other CBSA Dummy	−6.557***	−6.428***	−6.443***	−6.457***	−6.474***	−6.326***
	(0.036)	(0.035)	(0.036)	(0.038)	(0.037)	(0.037)
Same CBSA × Log Distance	−0.281***	−0.280***	−0.273***	−0.272***	−0.273***	−0.271***
	(0.003)	(0.003)	(0.004)	(0.004)	(0.004)	(0.004)
Other CBSA × Log Distance	−1.458***	−1.352***	−1.378***	−1.356***	−1.370***	−1.342***
	(0.035)	(0.034)	(0.033)	(0.038)	(0.036)	(0.039)
Other CBSA × ι^l		−0.011***	−0.011***	−0.011***	−0.011***	−0.018***
		(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Same CBSA × Log Dist × ι^r		−0.0002***	−0.0003***	−0.0003***	−0.0003***	−0.0002***
		(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)
Other CBSA × Log Dist × ι^l		−0.009***	−0.009***	−0.009***	−0.008***	−0.008***
		(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Same CBSA × Log Dist × Orig Log Median Income			−0.045***	−0.050***	−0.041***	−0.041***
			(0.006)	(0.007)	(0.007)	(0.007)
Other CBSA × Log Dist × Orig Log Median Income			0.364***	0.465***	0.394***	0.413***
			(0.049)	(0.074)	(0.073)	(0.076)
Same CBSA × Log Dist × Orig Child Share				0.090***	0.078***	0.087***
				(0.026)	(0.026)	(0.027)
Other CBSA × Log Dist × Orig Child Share				−0.093	−0.046	0.096
				(0.355)	(0.339)	(0.333)
Same CBSA × Log Dist × Orig Log Median Age				0.005	−0.003	0.003
				(0.026)	(0.025)	(0.027)
Other CBSA × Log Dist × Orig Log Median Age				−0.608**	−0.618**	−0.441*
				(0.258)	(0.246)	(0.249)
Other CBSA × Log HH Income					0.183***	0.194***
					(0.015)	(0.016)
Other CBSA × Age					−0.007	0.052***
					(0.018)	(0.020)
Same CBSA × Log Dist × Log HH Income					−0.003**	−0.004***
					(0.001)	(0.001)
Other CBSA × Log Dist × Log HH Income					0.107***	0.100***
					(0.012)	(0.017)
Same CBSA × Log Dist × Log Age					−0.003**	−0.004**
					(0.001)	(0.001)
Other CBSA × Dist × Log Age					0.032**	−0.015
					(0.014)	(0.017)
Orig ZCTA - Year FE	Yes	Yes	Yes	Yes	Yes	No
Dest ZCTA - Year FE	Yes	Yes	Yes	Yes	Yes	No
Orig ZCTA - Year - High (Income/Age) - Long (CBSA/Res) FE	No	No	No	No	No	Yes
Dest ZCTA - Year - High (Income/Age) - Long (CBSA/Res) FE	No	No	No	No	No	Yes
Individual Level Controls	No	No	No	No	Yes	Yes
Clustered SE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,533,398	1,533,398	1,533,398	1,533,398	850,593	850,593

Note: Components of regressors are: log 1+distance in meters, demeaned separately for within and cross-CBSA movers; years of tenure in the home ι^r and CBSA ι^l; demeaned origin neighborhood 2010 median log HH income, child share, and median age; Infutor reported demeaned log HH income and log age. The regressions only uses movers. Standard errors are clustered by destination neighborhood. In Column (6), fixed effects are interacted with indicator variables for high vs. low income, age, and long vs. short CBSA and residence tenure, based on median splits of each respective variable.

A7.[13]

6.2 Fixed Moving Costs

Fixed moving cost estimates are derived from individual level regressions that use Equation (17) to form the dependent variable, us explained in Section 5.4.

Depending on specification, the constant is between 0.732 and 0.931, implies an average fixed cost of about $40,000 for those with 0 years of tenure in the home. This estimate is quite robust to the details of how it is calculated and to the set of controls included. Whether inferring the value of staying for $\iota^r = 0$ out of sample or using actual ι^r and controlling for its difference between the two values of staying, estimates are very similar.

Fixed costs of moving decline markedly with income and increase with age. We find evidence of this using either neighborhood characteristics or Infutor reported characteristics. However, the estimated importance of demographics using the Infutor characteristics are much more muted. Based on neighborhood characteristics, increasing household income by 10 percent reduces fixed costs by $2,500 and increasing age by 10 percent reduces fixed cost by almost twice as much.

6.3 Neighborhood Substitution Patterns

Results of estimating the individual predicted probabilities gravity regression in equation (14) identify summary substitution parameters $\sigma_l^s - 1$ and $\sigma_l^o - 1$. Results show that those migrating out of low income neighborhoods find alternatives within the same CBSA to be close substitutes. This reflects an observed randomness with which people tend to move between low income neighborhoods. However, those moving out of young neighborhoods to other CBSAs tend to have low substitution elasticities across neighborhoods. Other interaction effects, including CBSA population, are small. These results are similar to their nested logit counterparts reported in Table A7.

6.4 Marginal Values and Utilities of Rents

Table 6 Columns (1)-(3) report individual level regressions of $\bar{V}_t^{stay}(l, n, \iota^l, \iota^r, Z_{it}) + \tau_0^s(\iota^r, Z_{it})$ on tenure lengths, neighborhood log rent, neighborhood fixed effects, and CBSA-year fixed effects. Results in remaining columns have flow utility plus re-scaled fixed moving cost as the outcome instead, derived in equation (16).

The first and third columns show results without instrumenting for rents. In these specifications, the rent coefficient is near 0, reflecting the fact that positive local amenity

[13] Interpretation of the PPML estimates requires rescaling by $\gamma\sigma$ rather than just γ, with σ estimates also reported in Table A7.

Table 5: Fixed Moving Cost Estimates

	Dependent variable:					
	$\tau_0^s + V_{it}^{stay}(l,n,\iota^l,\iota^r,Z) - V_{it}^{stay}(l,n,\iota^l,0,Z)$			$\tau_0^s + \hat{V}_{it}^{stay}(l,n,\iota^l,0,Z) - V_{it}^{stay}(l,n,\iota^l,0,Z)$		
	(1)	(2)	(3)	(4)	(5)	(6)
Res. Stay Year (ι^r)	0.143***	0.143***	0.143***			
	(0.0003)	(0.0002)	(0.0002)			
Log ZCTA Median Income		−0.453***			−0.373***	
		(0.028)			(0.025)	
Log ZCTA Median Age		−0.841***			−0.713***	
		(0.073)			(0.069)	
Log HH Income			−0.070***			−0.047***
			(0.007)			(0.006)
Log Age			−0.062***			−0.047***
			(0.006)			(0.005)
Constant	0.827***	0.905***	0.732***	0.869***	0.931***	0.793***
	(0.014)	(0.010)	(0.012)	(0.012)	(0.010)	(0.011)
Clustered SE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,533,398	1,533,398	850,593	1,533,398	1,533,398	850,593

Note: This table shows regression results comparing the value of staying versus moving to the same neighborhood. Columns (1)–(3) use estimates of stay utility based on observed characteristics and actual residential tenure. Columns (4)–(6) report analogous estimates, but with residential tenure set to zero. Age and income variables are demeaned. Standard errors are clustered at the neighborhood level.

shocks are likely capitalized into higher current and expected future rents. Instrumenting for rents using the neighborhood supply elasticity interacted with average Texas log rent, as explained in Section 3.2, yields an expected reduction in this coefficient, to -1.43 for values and -0.38 for utilities.

To think about the comparison of the value and utility rent elasticities, consider the following decomposition:

$$\frac{\mathrm{d}\bar{V}_{it}^{stay}(l,n)}{\mathrm{d}\ln Rent_{nt}} = \frac{\mathrm{d}u_{it}(l,n)}{\mathrm{d}\ln Rent_{nt}} + \beta \sum_{l',n'} p_{it+1}^{nn'} \frac{\mathrm{d}\bar{V}_{it+1}^{stay}(l',n')}{\mathrm{d}\ln Rent_{n't+1}} \mathbb{E}_t \frac{\mathrm{d}\ln Rent_{n't+1}}{\mathrm{d}\ln Rent_{nt}} \tag{19}$$

This estimate of -1.43 can be compared against the following benchmark. Suppose that no migration is possible and individuals expect any year t shock in log rent to be permanent (that is, rents follow a martingale). In this case, the marginal value would be $\frac{\alpha}{1-\beta}$, or 10 times α given a β of 0.9. In practice, the AR(1) coefficient within neighborhood for log rent is 0.8 in our data, meaning that with no migration one would rationally expect a ratio of 8. Instead, the ratio of coefficients is 3.7, with the difference reflecting the migration option. The marginal utility of rent is less negative for higher income people and for older people

Table 6: Neural Net Utility Parameters Estimation

	Dependent variable:					
	$V_{it}^{Stay} + \tau_0^s$			$u_{it} + (1-\beta)\tau_0^s$		
	(1)	(2)	(3)	(4)	(5)	(6)
Res. Stay Year (ι^r)	0.145***	0.145***	0.144***	0.029***	0.029***	0.029***
	(0.0001)	(0.0001)	(0.0001)	(0.0002)	(0.0002)	(0.0002)
CBSA Stay Year (ι^l)	0.003***	0.003***	0.003***	−0.002***	−0.002***	−0.001***
	(0.00004)	(0.00004)	(0.00004)	(0.00002)	(0.00003)	(0.00002)
Log Rent	−0.056*	−1.427***	−0.917***	−0.013	−0.382***	−0.280***
	(0.030)	(0.406)	(0.291)	(0.011)	(0.125)	(0.105)
Log Rent × Log Income			0.023			0.025**
			(0.030)			(0.011)
Log Rent × Log Age			−0.069			0.087**
			(0.055)			(0.038)
Rent IV	No	Yes	Yes	No	Yes	Yes
ZCTA Group FE	Yes	Yes	Yes	Yes	Yes	Yes
CBSA-Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Clustered SE	Yes	Yes	Yes	Yes	Yes	Yes
Other Controls	No	No	Yes	No	No	Yes
Observations	14,226,675	13,396,479	6,158,379	14,226,675	13,396,479	6,158,379

Note: This table presents regressions of the value of stay and flow utility from staying in a neighborhood on the log rent index. All interaction variables are demeaned. Other controls include individual level demeaned log income and log age. The regressions are run on the full sample of stayers and movers. Standard errors are clustered at the ZCTA group level.

(Column 6).

Estimated marginal utility of tenures in the home and neighborhood, ι^r and ι^l, are 0.029 and -0.001, respectively. For ι^r, the coefficient also incorporates the impact on the fixed cost of moving (which we could not estimate above). But comparison with the corresponding estimated log rent response in the $V_{it}^{stay} + \tau_0^s$ regression is instructive. This coefficient is 0.144 (Column 3). If that were fully accounted in the fixed cost, the coefficient in Column 6 would be 0.014. Instead, it is double that. Therefore, it must be that tenure in the home contributes to flow utility. With a ratio of 3.7 between the value and flow ι^r coefficients, we can solve out that marginal impact of ι^r on the fixed cost of moving is 0.05 ($2,500 per year) and on utility is 0.025. This means that the typical individual would have to be compensated $1,250 in lower rent to make her indifferent to losing one year of home tenure.

Table A8 presents corresponding nested logit-PPML based results of utility parameter estimation. This approach uses aggregate data only in a discretized state space. Nonetheless, rent results are similar with an estimated log rent coefficient of -0.40, which is not statistically different from -0.28. One advantage of the aggregate approach is that it provides an alternative and more quasi-experimental way to identify γ, jointly with utility parameters by

estimation equation (28). We can address the potential endogeneity of $IV^m_{lnt}(\iota^l, \iota^r, Z_{it})$ with an instrument based on an aggregate of the values of neighborhoods outside of ln measured in period $t+1$. Estimated in this way, we find γ=0.4 (Table A8).

6.5 Model Fit

The model is specified to flexibly fit dynamic patterns of migration across cities and neighborhoods. In the third estimation step, the neural net successfully fits the declining rates of out-migration seen in the data. However, the model is not able to distinguish whether these declining migration rates are because of growing fixed migration costs or growing local ties. We also note that as it is much more flexible, the neural net specification fits the data better in various ways than the more restricted aggregated nested logit specification. (Detailed statistics to be filled in here.) Figure A6 demonstrates that the first estimation step neural net successfully predicts migration counts in the hold-out test data set. A regression of out-of-sample predicted on actual movers yields an R-squared of 0.92 and a regression coefficient of 0.87.

7 The Welfare Burden of Rent Growth

We have estimated that the semi-elasticity of the value of staying with respect to rent is about -1. This number slightly overstates the true welfare burden of rent growth because it ignores the immediate move option encompassed in the full value. Instead, it considers a renter who is locked into renting a unit for at least one more year after rents have gone up. Using our estimates of the values of staying and moving plus moving costs, we will calculate compensating variations of actual rent growth for people in different demographic groups in different types of location in 2010. Then, we will formally investigate the extent to which these calculated welfare burdens have been due to migration costs, the potential lack of substitutibility across neighborhoods, and spatial correlations in rent growth.

We emphasize that simple knowledge of $\alpha(Z_{it})$ is not sufficient to measure the welfare burden of rent growth. The reason is that neighborhoods differ on their accessibility to outside options. Full differentiation of the value of staying in equation (**??**) reveals that migration probabilities to each potential destination times changes in continuation values in each possible location matters. So individuals for whom migration costs are low, thereby leading to higher migration probabilities, face lower burdens from local shocks.

Based on parameter estimates, preliminary evidence indicates that those in 2010 low rent neighborhoods in large CBSAs were most exposed to rising rents, with the largest welfare losses. But this is because other CBSAs were both very costly to move to and poor

substitutes for Dallas and Houston. Indeed, the cost of moving across CBSAs within Texas is estimated to be slightly greater than the fixed cost of moving, of about $40,000. Most of these welfare losses come from the lack of remaining low rent neighborhoods in the large CBSAs.

In a second set of counterfactuals, we will assess the welfare consequences of exogenously adding 1000 units of subsidized rental housing to each neighborhood, similar to the Low Income Housing Tax Credit (LIHTC) program, one at a time. This will reveal the types of neighborhoods that would benefit the most from project based subsidized housing.

8 Conclusions

This paper specifies and estimates a dynamic model of demand for neighborhoods by renters. This model is specified so as to focus on quantifying the burdens of rent growth. These burdens can come in the form of high moving costs to alternative affordable neighborhoods, the availability of poor substitute neighborhoods, and high correlations across neighborhoods in rent growth.

We estimate fixed moving costs of about $40,000 that are sharply decreasing in inferred household income. Moreover, we find even larger variable costs of moving between CBSAs. The typical variable cost of moving within a CBSA is less than $1,000. The highest burden of rent growth appears to be for residents of low rent neighborhoods in the largest CBSAs. While other neighborhoods in their same CBSA are reasonably good demand substitutes, these residents face high moving costs to other CBSAs and few to no alternative neighborhoods with low rents during our study period. These burdens are lower for younger individuals but higher for lower income individuals.

References

Almagro, M. and Domınguez-Iino, T. (2025), 'Location sorting and endogenous amenities: Evidence from amsterdam', *Econometrica* .

Ambrose, B. W., Coulson, N. E. and Yoshida, J. (2015), 'The repeat rent index', *Review of Economics and Statistics* **97**(5), 939–950.

Arcidiacono, P. and Miller, R. A. (2011), 'Conditional choice probability estimation of dynamic discrete choice models with unobserved heterogeneity', *Econometrica* **79**(6), 1823–1867.

Artuç, E. and McLaren, J. (2015), 'Trade policy and wage inequality: A structural analysis

with occupational and sectoral mobility', *Journal of International Economics* **97**(2), 278–294.

Baum-Snow, N. and Han, L. (2024), 'The microgeography of housing supply', *Journal of Political Economy* .

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

Bodéré, P. (2023), 'Dynamic spatial competition in early education: An equilibrium analysis of the preschool market in pennsylvania', *Job Market Paper* .

Caliendo, L., Dvorkin, M. and Parro, F. (2019), 'Trade and labor market dynamics: General equilibrium analysis of the china trade shock', *Econometrica* **87**(3), 741–835.

Correia, S., Guimarães, P. and Zylkin, T. (2020), 'Fast poisson estimation with high-dimensional fixed effects', *The Stata Journal* **20**(1), 95–115.

Couture, V. and Handbury, J. (2020), 'Urban revival in america', *Journal of Urban Economics* **119**, 103267.

Davis, M. A., Gregory, J., Hartley, D. A. and Tan, K. T. (2021), 'Neighborhood effects and housing vouchers', *Quantitative Economics* **12**(4), 1307–1346.

Diamond, R., McQuade, T. and Qian, F. (2018), 'The effects of rent control expansion on tenants, landlords, and inequality: Evidence from san francisco', *NBER Working Paper No. 24181* .

Diamond, R., McQuade, T. and Qian, F. (2019), 'The effects of rent control expansion on tenants, landlords, and inequality: Evidence from san francisco', *American Economic Review* **109**(9), 3365–94.

Dingel, J. I. and Tintelnot, F. (2020), Spatial economics for granular settings, Technical report, National Bureau of Economic Research.

French, R., Gandhi, A. and Gilbert, V. (2023), Quantifying the welfare impacts of gentrification on incumbent renters, Technical report, Working Paper.

Ganong, P. and Shoag, D. (2017), 'Why has regional income convergence in the us declined?', *Journal of Urban Economics* **102**, 76–90.

Guimaraes, P., Figueirdo, O. and Woodward, D. (2003), 'A tractable approach to the firm location decision problem', *Review of Economics and Statistics* **85**(1), 201–204.

Gyourko, J., Mayer, C. and Sinai, T. (2013), 'Superstar cities', *American Economic Journal: Economic Policy* **5**(4), 167–99.

Hotz, V. J. and Miller, R. A. (1993), 'Conditional choice probabilities and the estimation of dynamic models', *The Review of Economic Studies* **60**(3), 497–529.

Howard, G. and Shao, H. (2023), The dynamics of internal migration: A new fact and its implications, Technical report, Working paper.

Imbert, C., Monras, J., Seror, M. and Zylberberg, Y. (2023), Floating population: migration with (out) family and the spatial distribution of economic activity, Federal Reserve Bank of San Francisco.

Kalouptsidi, M., Scott, P. T. and Souza-Rodrigues, E. (2021), 'Linear iv regression estimators for structural dynamic discrete choice models', *Journal of Econometrics* **222**(1), 778–804.

Lin, W. and Wooldridge, J. M. (2019), Testing and correcting for endogeneity in nonlinear unobserved effects models, *in* 'Panel data econometrics', Elsevier, pp. 21–43.

Molloy, R. (2020), 'The effect of housing supply regulation on housing affordability: A review', *Regional science and urban economics* **80**(C).

Molloy, R., Nathanson, C. G. and Paciorek, A. (2022), 'Housing supply and affordability: Evidence from rents, housing consumption and household location', *Journal of Urban Economics* **129**, 103427.

Pakes, A., Porter, J. R., Shepard, M. and Calder-Wang, S. (2021), Unobserved heterogeneity, state dependence, and health plan choices, Technical report, National Bureau of Economic Research.

Ruggles, S., Flood, S., Goeken, R., Grover, J., Meyer, E., Pacas, J. and Sobek, M. (2020), 'Ipums usa: version 10.0 [dataset]', *Minneapolis, Mn: Ipums* **10**, D010.

Schubert, G. (2021), House price contagion and us city migration networks, Technical report, Working Paper.

Scott, P. (2014), 'Dynamic discrete choice estimation of agricultural land use'.

Train, K. E. (2009), *Discrete choice methods with simulation*, Cambridge university press.

A Nested Logit Model Specification and Estimation

This appendix describes the nested logit model specification we use as an alternative to the neural net. The nested logit specification uses the nesting structure of the model to discipline the predicted probabilities of migration choices. The neural net specification, described in the main text, instead uses flexible parameterization of value functions with observable characteristics of individuals and their neighborhoods as predictive features.

A.1 Nested Logit Model Specification

In the top nest, individuals decide to stay or move out of their home. In the middle level, they choose their commuting zone. In the lower level, they choose their neighborhood within the chosen commuting zone. These middle and lower level nests only apply conditional on moving. Mechanically, this specification resembles that in Imbert et al.'s (2023) investigation of rural-urban family migration in China. Each CBSA (nest) has its own dispersion parameter σ_l, which we allow to depend parametrically on CBSA attributes and whether the move is within or between CBSAs.

The nested logit model structure can be visualized in Figure A5. The joint distribution of ϵ_{ti} across staying in the home and moving to each possible neighborhood for individual i in CBSA l in year t has following cumulative distribution function:

$$\exp\left(-e^{-\epsilon^s_{ti}} - \left(\left(\sum_{n\in l} e^{-\epsilon^m_{lnti}/\gamma\sigma^s_l}\right)^{\sigma^s_l} + \sum_{l'\neq l}\left(\sum_{n\in l'} e^{-\epsilon^m_{l'nti}/\gamma\sigma^o_{l'}}\right)^{\sigma^o_{l'}}\right)^{\gamma}\right).$$

As with the neural net model specification, the upper nest is the move-stay decision. Similar to the neural net specification, the parameter γ describes the correlation between the composite inclusive values of each CBSA. As γ rises, the correlation declines, thereby making moving more attractive overall. Therefore, γ can be interpreted of as a general "preference" for moving. Alternatively and equivalently, γ can be viewed as a push factor out of one's home. This could be the rate of exogenous separation due to landlords' non-renewal of leases or the rate at which family structure changes, prompting idiosyncratically low values of staying.

This structure allows dispersion parameters, and hence elasticities of substitution across neighborhoods, to differ across commuting zones conditional on observables. In the neural net specification differences in observables across locations instead drive such variable substitution elasticities. The correlation between errors in commuting zone l' is increasing in $1-\sigma_{l'}$, where the restriction $0 \leq \sigma_{l'} \leq 1$ and $0 \leq \gamma \leq 1$ ensures that preferences are well-behaved (Train, 2009). We additionally index σ_l by whether it is associated with moving

within the same CBSA (subscript s) or to another CBSA (subscript o). This is intended to capture the idea that individuals may have better information about different neighborhood options in their current CBSA than in other CBSAs. In estimation, we allow the dispersion parameters σ_l^s and σ_l^o to depend on commuting zone population. The goal is to capture the fact that neighborhoods may have different substitution elasticities in different types of locations.

Given the nested logit specification of shocks, we can write the movers' ex-ante value function $\bar{V}_t^{move}(l, n, \iota^l, \iota^r) = \mathbb{E}_t V_t^{move}(l, n, \iota^l, \iota^r)$ in equation (9) as:

$$\bar{V}_t^{move}(l, n, \iota^l, \iota^r) = V_0^{move}(\iota^r) + \gamma \ln \left[\left(\sum_{n' \in N_l} \exp \frac{\bar{V}_t^{stay}(l, n', \iota^l, 0) - \tau_d^s(\iota^r) \ln d_{nn'}}{\gamma \sigma_l^s} \right)^{\sigma_l^s} + \sum_{l'=1, \neq l}^{L} \left(\sum_{n' \in N_{l'}} \exp \frac{\bar{V}_t^{stay}(l', n', 0, 0) - \tau_0^o(\iota^l) - \tau_d^o(\iota^l) \ln d_{nn'}}{\gamma \sigma_{l'}^o} \right)^{\sigma_{l'}^o} \right] \tag{20}$$

The second term in this expression is written as $\gamma IV^m(l, n, \iota^l, \iota^r)$ in the main text. This is γ times the inclusive value of moving. After moving, agents receive a "renewal value" that reflects the fact that tenure in the home ι^r renews to 0. Tenure in the CBSA also renews to 0 if moving across CBSAs. In addition, the inclusive value of moving incorporates the moving cost $\tau_{nn'}(\iota^l, \iota^r)$ to each neighborhood and dispersion parameters.

A.2 Conditional Migration Probabilities

Without loss of generality, decompose the post-migration value of staying into a constant, a CBSA-specific mean, and a neighborhood-specific component:

$$\bar{V}_t^{stay}(l, n, \iota^l, 0) = V_{0t}^s + W_{lt}(\iota^l, 0) + Y_t(l, n, \iota^l, 0) \tag{21}$$

For representing associated components of continuation values and probabilities, it is convenient to define the following inclusive values, which themselves are components of the inclusive value of moving in equation (20):

$$IV_{lnl't}(\iota^l, \iota^r) = \begin{cases} \ln \sum_{n' \in l'} \exp \left[(Y_t(l, n', \iota^{l'}, 0) - \tau_d^s(\iota^r) \ln d_{nn'}) / \gamma \sigma_l^s \right] & \text{if } l' = l \\ \ln \sum_{n' \in l'} \exp \left[(Y_t(l', n', 0, 0) - \tau_d^o(\iota^l) \ln d_{nn'}) / \gamma \sigma_{l'}^o \right] & \text{if } l' \neq l. \end{cases} \tag{22}$$

The top line denotes the inclusive value of commuting zone l at time t from the perspective of individuals living in neighborhood n of the same commuting zone. The bottom line summarizes the value of commuting zone l' at time t from the perspective of individuals

living in a different commuting zone. These inclusive values depend on origin neighborhood because of origin neighborhood specific variation in migration costs. They do not depend on mean destination CBSA values $W_{lt}(\iota^l, 0)$ or fixed migration costs, as these factor out of the summations in equation (20). Therefore, the inclusive value of moving can be rewritten as

$$IV^m_{lnt}(\iota^l, \iota^r) = \ln \sum_{l'} \exp[(W_{l't}(\iota^{l'}, 0) - \mathbf{1}(l' \neq l)\tau^o_0(\iota^l))/\gamma + \sigma_{l'} IV_{lnl't}(\iota^l, \iota^r)].$$

Notice that the fixed migration cost $\tau^s_0(\iota^r)$ has been factored out.

The following probabilities are observed in the data conditional on living in a particular home in commuting zone l for ι^r and ι^l periods of time, respectively. Once conditioning on these objects, these probabilities do not differ across individuals i (assuming no imposition of individual-specific heterogeneity). Each probability listed below represents the fraction of migrants with residence and commuting zone tenures ι^r and ι^l that make each indicated location decision. These probabilities will be used in estimation of the nested logit specification of the model.

Conditional on moving, the probability of moving from neighborhood n to n' within commuting zone l is

$$p_t(l, n') = \frac{\exp\left((Y_t(l, n', \iota^{l'}, 0) - \tau^s_d(\iota^r) \ln d_{nn'})/\gamma\sigma^s_l\right)}{\exp\left(IV_{lnlt}(\iota^l, \iota^r)\right)} \frac{\exp\left(W_{lt}(\iota^{l'}, 0)/\gamma + \sigma^s_l IV_{lnlt}(\iota^l, \iota^r)\right)}{\exp\left(IV^m_{lnt}(\iota^l, \iota^r)\right)} \tag{23}$$

This is the probability of choosing neighborhood n' conditional on remaining in CBSA l times the probability of choosing to remain in CBSA l.

The probability of moving from neighborhood n in commuting zone l to neighborhood n' in commuting zone l' is

$$p_t(l', n') = \frac{\exp\left((Y_t(l', n', 0, 0) - \tau^o_d(\iota^l) \ln d_{nn'})/\gamma\sigma^o_{l'}\right)}{\exp\left(IV_{lnl't}(\iota^l, \iota^r)\right)} \frac{\exp\left((W_{l't}(0, 0) - \tau^o_0(\iota^l))/\gamma + \sigma^o_{l'} IV_{lnl't}(\iota^l, \iota^r)\right.}{\exp\left(IV^m_{lnt}(\iota^l, \iota^r)\right)} \tag{24}$$

In the next section we show how we identify each of these objects in estimation.

A.3 Estimation Methodology

The first two steps of estimation of the neural net specification in the main text reduce to one iterated high dimensional fixed effect PPML gravity regression in the nested logit specification. In order to estimate this version of the model, we need to discretize the state space so that there are enough individuals of each origin type prior to migration that this origin state can essentially be controlled for with origin fixed effects. We achieve this by segmenting tenure in the CBSA and home into "short" and "long" periods of 0-3 years

and greater than 3 years, respectively. The use of PPML and maximum likelihood has the advantage of obviating the common process of smoothing conditional choice probabilities across the type and state space in a first step. The nested logit model aggregates to a Poisson model, which can be feasibly estimated with high dimensional fixed effects.

The likelihood function reflects the probabilities that each agent i residing in location $\{l, n\}$ transitions to location $\{l', n'\}$ at the end of period t. In addition to dependence on prior location, probabilities may also depend on commuting zone and residence tenures, ι^l and ι^r, and model parameters. Using the expressions for probabilities laid out in the prior sub-section, the likelihood contribution of individual i at time t can be expressed as follows. In (25), $\mathbf{1}_{it}(l', n')$ is an indicator that equals 1 if individual i migrates to neighborhood $\{l', n'\}$ in year t. All probabilities are as in (23) and (24) above.

$$\mathcal{L} = \prod_i \prod_t \prod_{l'} \prod_{n' \in l'} p_t(l', n' | l, n, \iota^l, \iota^r; \Theta)^{\mathbf{1}_{it}(l', n')} \tag{25}$$

As flow utilities depend on year, all values and hence probabilities are time indexed. However, moving costs are not time indexed. Therefore, estimation is more efficient pooling all years together. However, the Infutor data set is too large to estimate this model using individual level data by maximum likelihood. Feasible estimation requires inferring parameters through aggregate migration flows across neighborhoods. Grouping the probabilities listed above across individuals, the log likelihood can be written as

$$\ln \mathcal{L} = \sum_t \sum_{\iota^l} \sum_{\iota^r} \sum_l \sum_n \sum_{l'} \sum_{n'} N_t(l', n' | l, n, \iota^l, \iota^r) \ln p_t(l', n' | l, n, \iota^l, \iota^r). \tag{26}$$

Note that the first four summations are over states whereas the next two are over controlled choices and the last one is over time.[14]

Guimaraes et al. (2003) shows that this aggregated conditional logit model is equivalent to a Poisson model (up to a known constant) and can be estimated by Poisson pseudo maximum likelihood (PPML). In the language of Guimaraes et al. (2003), individuals with the same state vector are in the same "group", and group fixed effects must be included as a control. Dingel and Tintelnot (2020) note the advantages of estimation using PPML in spatial granular settings such as ours.

[14] This aggregation ensures that that all person-years receive equal weight in estimation, just as they would if estimated as a nested logit.

A.4 Estimation Stage 1 Using PPML: Conditional Values and Substitution Parameters

This version of the first stage of estimation uses information on migration within and between commuting zones in a gravity PPML regression to recover parameters in (26). The use of gross flows to recover choice-specific values using PPML in the first stage of dynamic discrete choice model estimation is also seen in Artuç and McLaren (2015). This is an appealing alternative to the conditional choice probabilities (CCP) approach (Hotz and Miller, 1993) in settings like ours that feature a large number of choices and many 0 flows as a result.

The estimated values and migration costs depend on the state variables ι^l and ι^r. So as to limit the state space of the model, we allow ι^l and ι^r to take on three values each. 0 is for a new arrival, 1 is for tenure of less than 3 years, and 2 is for tenure of at least 3 years. Neighborhood choice probabilities for those that remain in the same commuting zone, given by (23) and (24), depend on both current and one period ahead ι^l and ι^r. To reduce computational burden, we assume that $\iota^{l\prime} = \iota^l$ except if ι^l=0.

The estimation equation that corresponds to (26) is the following nonlinear gravity regression, which we estimate by PPML with high dimensional fixed effects (Correia et al., 2020) using an iterative process. This equation is derived from (23) and (24), with parameter transformations indicated with tildes and defined below.

$$\begin{aligned} N_t(l', n'|l, n, \iota^l, \iota^r) = \exp\Big[& \\ \mathbf{1}(l' = l)[\tilde{Y}_{ln't}(\iota^l, 0) + (\tilde{\tau}_{dn} + \tilde{\tau}_{\iota r}\iota^r)\ln d_{nn'} + \tilde{W}_{lt}(\iota^l, 0) + \tilde{\sigma}^s(X_l)IV_{lnt}(\iota^l, 0)] & \\ +\mathbf{1}(l' \neq l)[\tilde{Y}_{l'n't}(0,0) + (\tilde{\tau}_{dl} + \tilde{\tau}_{\iota l}\iota^l)\ln d_{nn'} + \tilde{W}_{l't}(0,0) + \tilde{\sigma}^o(X_{l'})IV_{l'n't}(0,0)] & \\ +\lambda_{lnt\iota^l\iota^r}\Big]\nu_{lnl'n'\iota^l\iota^r t} & \end{aligned} \tag{27}$$

We estimate the above equation using PPML regressions of movements from origin to destination neighborhoods for different values of ι^l and ι^r. To deal with nonlinearities in values, we estimate parameters using an iterative procedure. In the first iteration, the regression includes origin neighborhood-year-ι^l-ι^r fixed effects, destination neighborhood-year-ι^l fixed effects for movements within CBSAs, destination neighborhood-year fixed effects for movements between CBSAs, and log distance between origin and destination neighborhoods interacted with within versus cross-CBSA move indicators, ι^r, and ι^l. That is, destination fixed effects are separately estimated for within and across CBSA movements and for different values of ι^l given within-CBSA movements. Fixed migration costs are not separately identified in this step of estimation.

These initial estimates of destination fixed effects capture the parameter clusters $\tilde{Y}_{ln't}(\iota^l, 0)+\tilde{W}_{lt}(\iota^l, 0)+\tilde{\sigma}^s(X_l)IV_{lnt}(\iota^l, 0)$ and $\tilde{Y}_{l'n't}(0,0)+\tilde{W}_{l't}(0,0)+\tilde{\sigma}^o(X_{l'})IV_{l'n't}(0,0)$ in equation (27)

for within and between CBSA migration flows, respectively. To separate out their components, these initial fixed effect estimates are demeaned within CBSA to form $\tilde{Y}_{ln't}(\iota'', 0)$ and $\tilde{Y}_{l'n't}(0,0)$. These objects and variable migration costs $(\tilde{\tau}_{dn}+\tilde{\tau}_{\iota r}\iota^r)\ln d_{nn'}$ and $(\tilde{\tau}_{dl}+\tilde{\tau}_{\iota l}\iota^l)\ln d_{nn'}$ are accumulated into inclusive values $IV_{lnlt}(\iota^{l'}, 0)$ and $IV_{lnl't}(0,0)$ using equation (22). Then, the estimation is repeated with these two inclusive values included as additional regressors interacted with destination CBSA characteristics X_l. New estimates of $\tilde{Y}_{ln't}(\iota'', 0)$ and $\tilde{Y}_{l'n't}(0,0)$ using demeaned destination neighborhood fixed effects are calculated, and the process iterates until parameter convergence.

The structural interpretation of the "origin" fixed effects $\lambda_{lnt\iota^l\iota^r}$ is $\ln s_t(l, n, \iota^l, \iota^r) - IV^m_{lnt}(\iota^l, \iota^r)$, where $s_t(l, n, \iota^l, \iota^r)$ is the population of movers in each state (Guimaraes et al., 2003). This interpretation arises from the condition that probabilities of choosing all destination neighborhoods must sum to one for each combination of ι^l and ι^r. The second term in the "origin" fixed effect is the denominator of the choice probabilities in equations (23) and (24), which is the inclusive value of moving.

Each estimated parameter maps to corresponding structural parameters as follows. $\tilde{Y}_{ln't}(\iota'', 0) = Y_{ln't}(\iota'', 0)/\gamma\sigma^s_l$ is the adjusted value associated with moving to a different neighborhood in the same commuting zone. $\tilde{Y}_{l'n't}(0,0) = Y_{l'n't}(0,0)/\gamma\sigma^o_{l'}$ is the adjusted value associated with moving to a neighborhood in a different commuting zone. Adjusted variable migration cost parameters are, similarly, the structural parameters divided by $\gamma\sigma^s_l$ for within-CBSA migrants and $\gamma\sigma^o_{l'}$ for cross-CBSA migrants. Coefficients on inclusive values $\tilde{\sigma}^s(X_l)$ and $\tilde{\sigma}^o(X_{l'})$ are $\sigma^s_l(X_l) - 1$ and $\sigma^o_{l'}(X_{l'}) - 1$, respectively. These are the negative of correlations of the error term across neighborhoods within a commuting zone. A higher correlation implies that neighborhoods are more substitutable, and we expect this correlation to vary with the size of the commuting zones. We assume a linear relationship for the dispersion parameter as a function of destination CBSA population $(\tilde{\sigma^x_l} = \alpha^x + \beta^x \ln Pop_l, x \in \{s, o\})$.[15]

A.5 Estimation Stage 2 Using PPML: An Aggregate Euler Equation

The approach using individual level data has the advantage of leveraging maximum individual level heterogeneity to recover flow utility parameters. It also has the advantage of directly delivering estimates of $\bar{V}^{stay}_t(l, n, \iota^l, \iota^r, Z_{it})$ that are useful in welfare analysis. However, it has the disadvantage of less transparency in the identification of γ and some difficulty with standard error calculations in the flow utility decomposition.

As an alternative and complementary approach, we isolate flow utilities by differencing

[15] We confirm that estimates yield $\sigma^x_l > 0, \forall l, x \in \{s, o\}$. This condition is required for neighborhood choice probabilities to be increasing in their values.

out continuation values using aggregate neighborhood-year-ι level data on stay probabilities in years t and $t+1$. This approach has the disadvantage of accommodating less individual level heterogeneity. However, since most such heterogeneity is at the neighborhood level, we see this as a minor drawback. The following structural regression equation is derived below.

$$\begin{aligned}&\ln p^{stay}_{lnt}(\iota^l,\iota^r) - \ln p^{move}_{lnt}(\iota^l,\iota^r) + \beta \ln p^{move}_{lnt+1}(\iota^l,\iota^r) \\ &\quad = u_t(l,n,\iota^l,\iota^r) - (1-\beta)V^m_0(\iota^r) - \gamma[IV^m_t(l,n,\iota^l,\iota^r) - \beta IV^m_{t+1}(l,n,\iota^{l'},\iota^{r'})] + \xi_{lnt\iota^l\iota^r}\end{aligned} \tag{28}$$

The derivation of this equation uses the assumption that individuals have rational expectations about components of neighborhood values in year $t+1$. The difference between such expectations and realizations of $IV^m_{t+1}(l,n,\iota^{l'},\iota^{r'})$ leads to the regression error term ξ_t.

Here is the derivation:

$$\begin{aligned}&\ln(p^{stay}_t/p^{move}_t) = \ln\left(\frac{\exp V^{stay}_t}{\exp V^{move}_t}\right) = V^{stay}_t - V^{move}_t = (u_t + \beta E\bar{V}_{t+1}) - V^{move}_t \\ &\ln(p^{move}_{t+1}) = V^{move}_{t+1} - \bar{V}_{t+1} \\ &\implies \ln(p^{stay}_t/p^{move}_t) + \beta\ln(p^{move}_{t+1}) \equiv \ln p^{stay}_t - (\ln p^{move}_t - \beta \ln p^{move}_{t+1}) = \\ &u_t + \beta(E_t\bar{V}_{t+1} - \bar{V}_{t+1}) - (V^{move}_t - \beta V^{move}_{t+1}) = u_t - (V^{move}_t - \beta V^{move}_{t+1}) + \xi_t = \\ &u_t - (1-\beta)V_0 - \gamma_n(IV^{move}_t - \beta IV^{move}_{t+1}) + \xi_t\end{aligned}$$

Credible empirical implementation requires dealing with the potential endogeneity of both rents and the time differenced inclusive value of moving. On the latter issue, we can use $\gamma = 0.2$ estimated with the neural net and move the term including inclusive values to the left hand side. Utility function parameter estimates are insensitive to this choice of γ within a reasonable range. Alternatively, we can instrument for $IV^m_t(l,n,\iota^l,\iota^r) - \beta IV^m_{t+1}(l,n,\iota^{l'},\iota^{r'})$ with the component that only includes neighborhoods beyond some distance buffer of ln. We instrument for log rent as in Section 3.2.

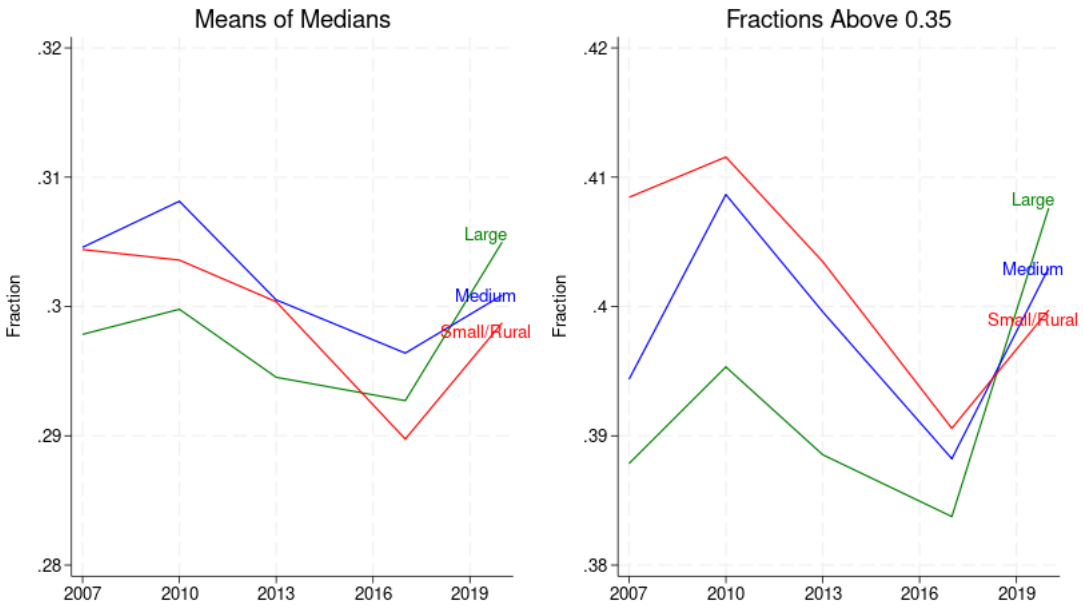
The finite dependence property of the model is key in the derivation of equation (28), as it allows us to use common renewal values across all agents. Unlike in Arcidiacono and Miller (2011) and Diamond et al. (2018) we use all of the data for this residualization, and not just comparisons across origin states in the periods prior to moving to renewal states. This has the advantage of giving us broader based identification of flow utilities, at the cost of using the nested logit distributional assumptions to construct closed forms for continuation values and the value of moving. Indeed, equation (28) can be viewed as a selection equation, in which the differenced inclusive value term controls for selection on migration.

The left-hand side of equation (28) is fully observed in the data, measured by the aggregate number of stayers and movers in each neighborhood-year for each tenure level, ι^l

and ι^r. Movers are individuals who move out of their residence and stayers are those that have the same addresses in both periods. On the right-hand side, we include the flow utility parameters from equation (4) and add the difference between the current and next period's inclusive value of moving, calculated in an earlier step of estimation, to identify the parameter γ. To account for differences in neighborhood amenities, we include ZCTA fixed effects. Additionally, we add CBSA-year fixed effects to capture variation in the valuation of CBSAs, such as changes in labor market conditions.

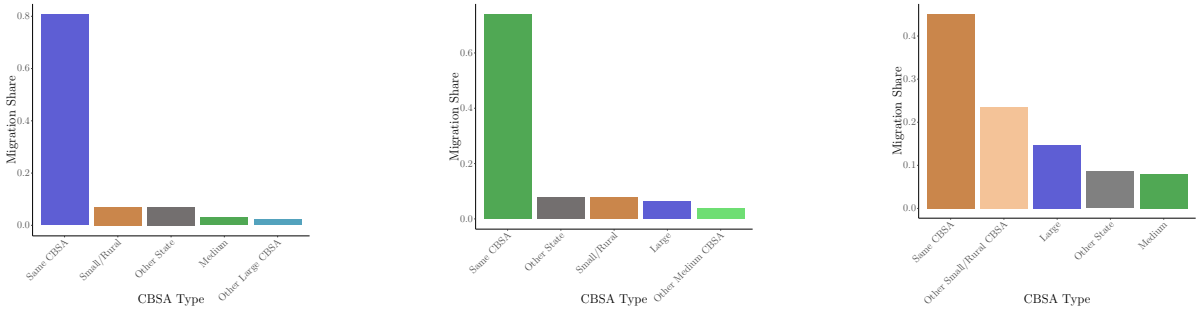
B Appendix Figures

Figure A1: Rent to Income Ratios, 2000-2020



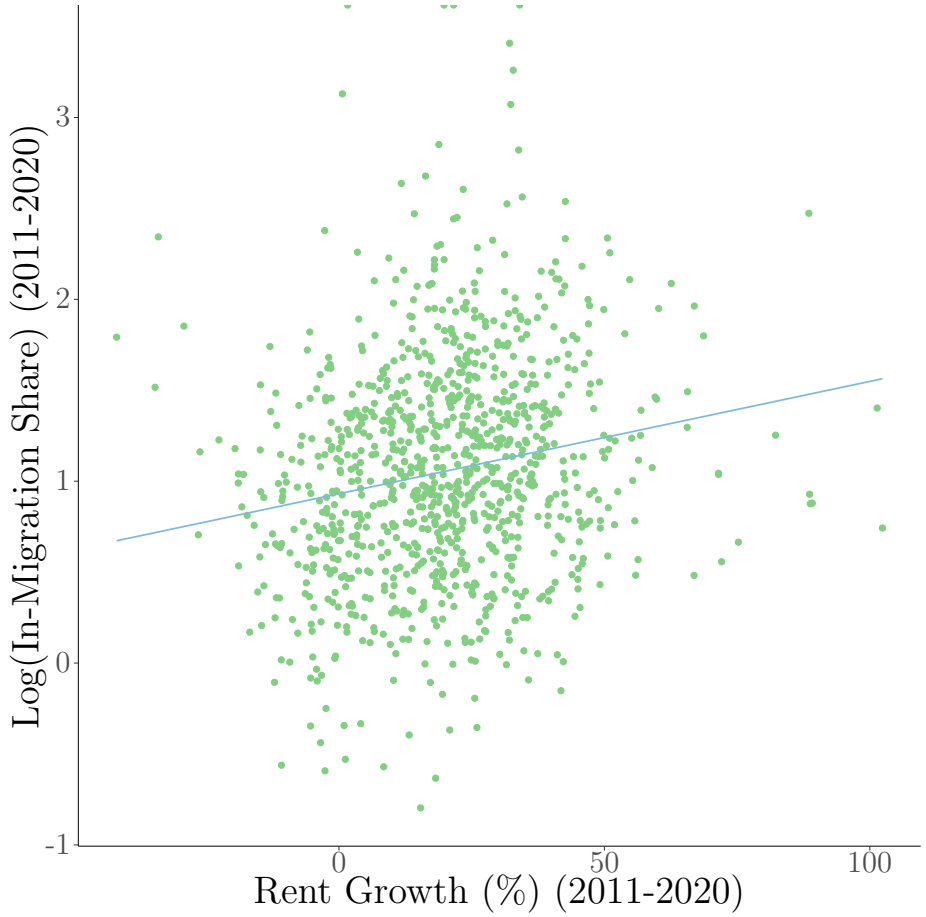
Notes: Each line in the left panel is calculated as the renter household weighted mean median rent to income ratio across neighborhoods in each indicated category. The right panel shows fractions of households paying at least 35% of household income in gross rent, averaged across neighborhoods.

Figure A2: Movement Across CBSA Sizes



(a) Large CBSAs (b) Medium CBSAs (c) Small CBSAs

Figure A3: In-Migration Rates and Rents



Notes: Plot shows the relationship between the change in the share of the neighborhood population that has arrived in the prior year and the growth in the Zillow rent index. Each neighborhood in our data gets one observation in each year 2011-2020.

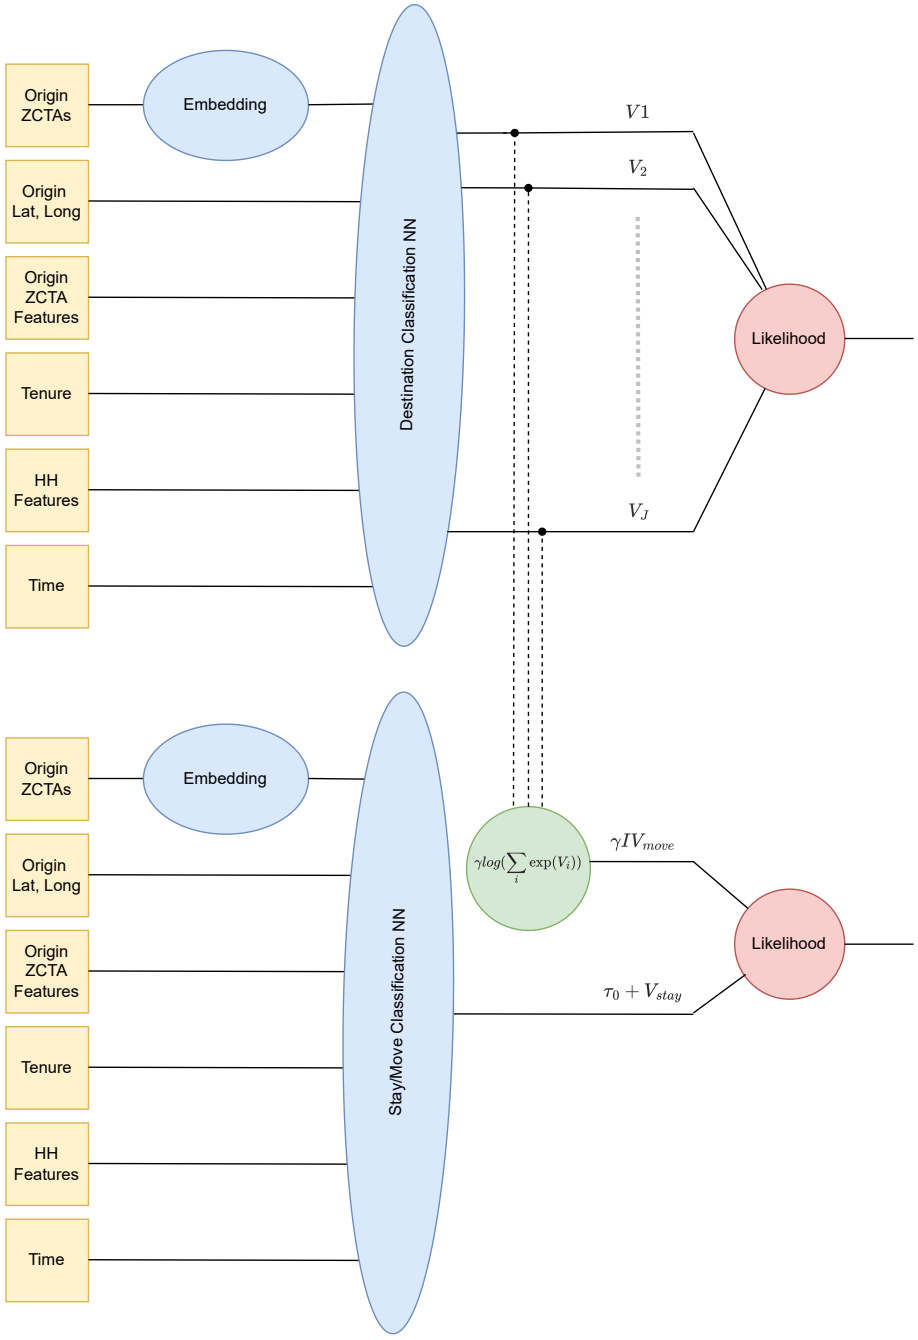
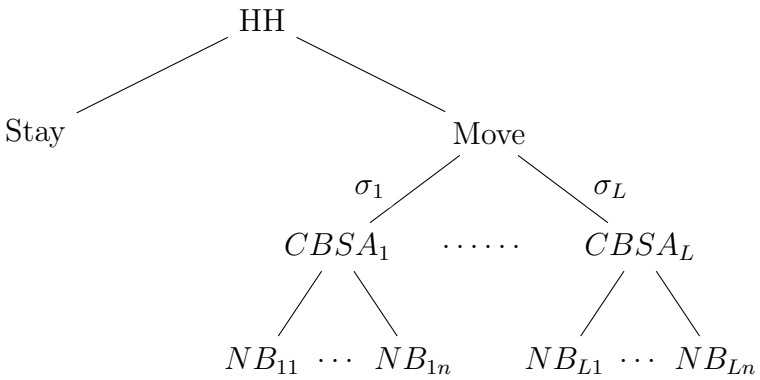


Figure A4: Nested Neural Network Model

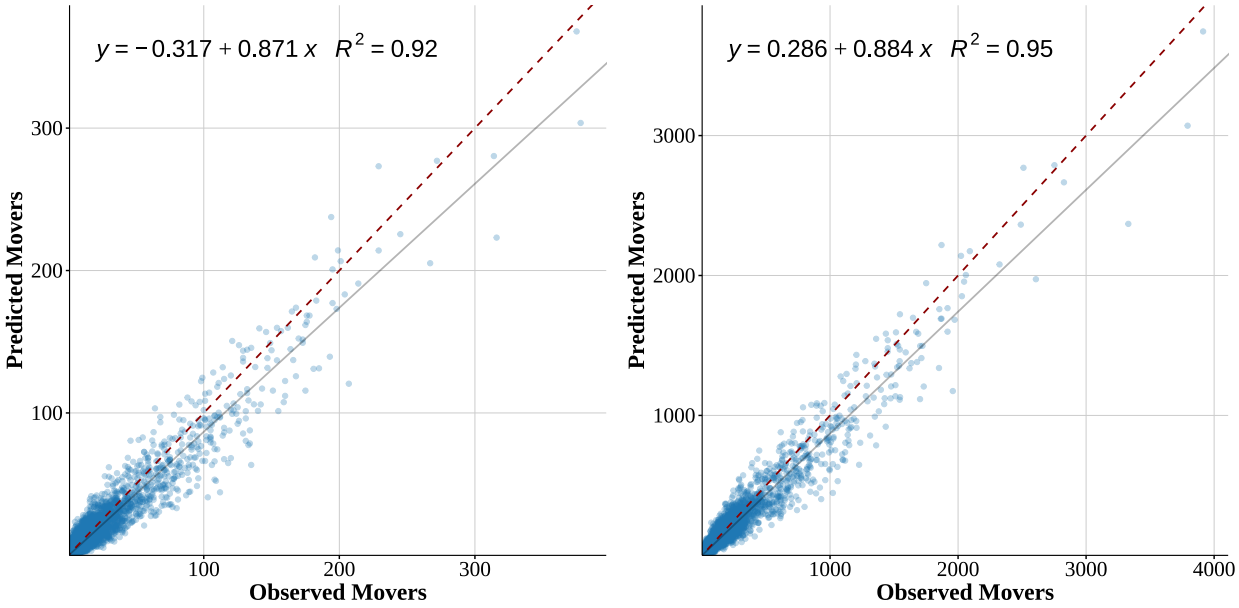
Note: This figure shows the nested architecture of the estimated neural network model. The destination classification subnetwork is trained on movers and predicts the value of moving to each destination ZCTA group. The stay/move classification subnetwork aggregates these destination utilities—conditional on moving, using a log-sum-exp operation scaled by a parameter γ, and then computes the relative value of staying. The subnetworks are trained separately by maximizing likelihood-based objective functions.

Figure A5: Nesting Structure in the Nested Logit Specification



Note: In the nested logit specification, each CBSA nest has its own dispersion parameter σ_l. The neural net specification has the dispersion parameter σ across all neighborhoods. Both specifications have the dispersion parameter γ associated with the shock to the value of staying.

Figure A6: Predicted vs Actual Movement Flows



(a) Test Subset (b) Entire Data

Note: Each point in the scatter plot shows the number of movers for a given origin–destination pair in a specific year.

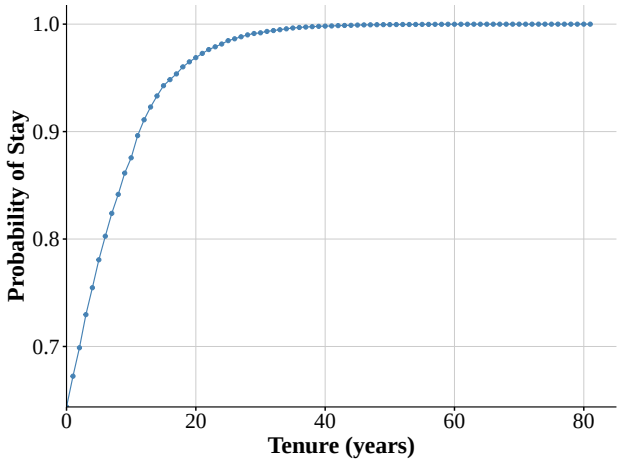


Figure A7: Tenure vs Probability of Stay in the Same Residence

C Appendix Tables

Table A1: Renter Stocks and Flows (millions)

			Moves			Stays
Year	First	Last	Into TX	Within	Out of TX	
2010	4.3	0.5	0.3	0.4	0.0	3.6
2011	0.3	0.3	0.3	0.3	0.0	3.6
2012	0.3	0.4	0.4	0.3	0.0	3.5
2013	0.3	0.3	0.3	0.3	0.0	3.5
2014	0.5	0.6	0.5	0.4	0.0	3.4
2015	0.4	0.4	0.4	0.4	0.0	3.4
2016	0.5	1.3	0.6	0.4	0.0	3.4
2017	0.4	0.6	0.4	0.4	0.0	2.6
2018	0.4	0.6	0.4	0.3	0.1	2.4
2019	0.4	2.8	0.4	0.3	0.1	2.1
Total	7.7	7.7	4.0	3.6	0.3	31.5

tNote: "First" indicates the number of individuals entering the data set in the indicated year. "Last" indicates the number leaving the data set in that year. The following three columns show the number of moves into, within, and out of Texas. The final column shows the number of individuals who remain in their home throughout the indicated calendar year.

Table A2: Owner Stocks and Flows (millions)

			Moves			Stays
Year	First	Last	Into TX	Within	Out of TX	
2010	11.8	0.7	0.7	0.9	0.0	10.4
2011	0.5	0.5	0.5	0.5	0.0	10.7
2012	0.7	0.6	0.7	0.6	0.0	10.7
2013	0.5	0.6	0.6	0.7	0.0	10.6
2014	1.0	1.2	1.0	0.7	0.0	10.6
2015	0.6	0.7	0.7	0.7	0.0	10.3
2016	0.9	3.5	1.0	0.8	0.0	10.2
2017	0.6	1.3	0.7	0.7	0.1	7.7
2018	0.6	1.4	0.7	0.6	0.1	7.0
2019	0.6	7.4	0.7	0.6	0.1	6.2
Total	18.0	18.0	7.2	6.8	0.5	94.4

Note: This table is analogous to Table (A1), except it uses Infutor information on owners instead of renters.

Table A3: Correlation between the Rent Instrument and ZCTA Group Demographics

	Dependent Variable:						
	Log ZORI	Log Rent	Log Income	Log Age	Log Rental Units	Children	Education
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
IV	−0.32***	−0.30***	−0.01	−0.02	0.11	0.03	−0.01
	(0.06)	(0.08)	(0.08)	(0.04)	(0.19)	(0.03)	(0.02)
ZCTA FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
CBSA - Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Clustered	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	8,121	8,209	8,210	8,210	8,210	8,210	8,210
R^2	0.95	0.96	0.98	0.95	0.98	0.96	0.98
Note:							*p<0.1; **p<0.05; ***p<0.01

Note: This table reports the correlation between the rent instrument and time-varying demographic characteristics at the ZCTA level. The dependent variables include the logarithm of the ZORI index, ACS median rent, median income, median age, total renter-occupied units, and the share of households with children and those with a bachelor's degree. All standard errors are clustered at the ZCTA group level.

Table A4: Static Model PPML Regressions for Movers and Stayers

Dependent Variable:	Number of Migrants			
	PPML-OLS		PPML-IV	
Model:	(1)	(2)	(3)	(4)
Variables				
Log Rent	-0.23***	-0.31***	-1.0*	-1.1*
	(0.05)	(0.07)	(0.59)	(0.60)
Mover Dummy	-8.2***	-2.5***	-2.5***	-2.6***
	(0.02)	(0.010)	(0.010)	(0.01)
Log Distance		-0.53***	-0.53***	
		(0.001)	(0.001)	
Stage 1 Residual			0.71	0.81
			(0.59)	(0.60)
Log Distance (Same CBSA)				-0.43***
				(0.002)
Log Distance (Other CBSA)				-0.62***
				(0.002)
Fixed-effects				
Destination ZCTA Group	Yes	Yes	Yes	Yes
Origin ZCTA Group - Time	Yes	Yes	Yes	Yes
Destination CBSA - Time	Yes	Yes	Yes	Yes
Fit statistics				
Observations	7,545,984	7,545,984	7,545,984	7,538,672
Pseudo R^2	0.89667	0.95643	0.95643	0.96358

Clustered (Destination ZCTA Group) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Note: The dependent variable is the number of migrants from the origin ZCTA to the destination ZCTA in each year, including those who do not move. Columns (1) and (2) present the correlations, while columns (3) and (4) show PPML w-IV regressions. Stage 1 Residual is the error of prediction from the first stage. "Distance" measures the distance between the origin and destination, and "Distance (Same CBSA)" is the distance interacted with a dummy variable equal to 1 if the move occurs within the same CBSA.

Table A5: PPML Regression with Heterogeneity in Moving Cost for Movers and Stayers

Dependent Variable:	Number of Migrants			
Model:	(1)	(2)	(3)	(4)
Variables				
Log Rent	-1.8***	-1.8***	-1.5***	-1.7***
	(0.34)	(0.34)	(0.31)	(0.29)
Mover Dummy	-3.7***	-3.7***	-3.7***	-3.7***
	(0.01)	(0.01)	(0.01)	(0.01)
Stage 1 Residual	1.6***	1.6***	1.3***	1.5***
	(0.34)	(0.34)	(0.31)	(0.29)
Log Distance (Same CBSA)	-0.43***	-0.43***	-0.35***	-0.36***
	(0.003)	(0.002)	(0.002)	(0.002)
Log Distance (Other CBSA)	-0.62***	-0.62***	-0.61***	-0.62***
	(0.003)	(0.002)	(0.002)	(0.002)
Log Distance * Low Income Share (Same CBSA)	-0.07***			-0.21***
	(0.02)			(0.02)
Log Distance * Low Income Share (Other CBSA)	-0.13***			-0.23***
	(0.02)			(0.009)
Log Distance * Young Share (Same CBSA)		0.30***		0.14***
		(0.03)		(0.02)
Log Distance * Young Share (Other CBSA)		0.24***		0.18***
		(0.03)		(0.01)
Log Distance * CBSA Population (Same CBSA)			-0.07***	-0.07***
			(0.001)	(0.001)
Log Distance * CBSA Population (Other CBSA)			-0.009***	-0.03***
			(0.001)	(0.002)
Fixed-effects				
Destination ZCTA Group	Yes	Yes	Yes	Yes
Origin ZCTA Group - Time	Yes	Yes	Yes	Yes
Destination CBSA - Time	Yes	Yes	Yes	Yes
Fit statistics				
Observations	7,531,360	7,531,360	7,353,092	7,345,960
Pseudo R^2	0.98731	0.98741	0.98805	0.98855

Clustered (Destination ZCTA Group) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Note: This table reports the PPML regressions. The dependent variable is the number of migrants from the origin to destination neighborhood in each year. "Distance" measures the distance between the origin and destination, and "Distance (Same CBSA)" is the distance interacted with a dummy variable equal to 1 if the move occurs within the same CBSA. "Low Income Share" is the share of renters with less than $35K. "Young Share" is the share of renters between 15 to 35 years old. "CBSA Population" the population of the destination CBSA.

Table A6: Stay Probability as a Function of Rent

	Dependent Variable:		
	log(Probability of Stay)		
	(1)	(2)	(3)
log(Rent)	0.01***	0.21***	0.12***
	(0.003)	(0.05)	(0.04)
IV	No	Yes	Yes
Controls	No	No	Yes
CBSA - Year FE	Yes	Yes	Yes
ZCTA FE	Yes	Yes	Yes
Clustered	Yes	Yes	Yes
Observations	9,112	7,298	7,312
Note:	*p<0.1; **p<0.05; ***p<0.01		

Table A7: Dynamic Model First Stage Regression

Dependent Variable:	Number of Migrants			
Model:	(1)	(2)	(3)	(4)
Variables				
Log Distance (Same CBSA)	-0.41***	-0.41***	-0.42***	-0.42***
	(0.005)	(0.005)	(0.004)	(0.004)
Log Distance (Same CBSA) * Long-term Residence	0.03***	0.03***	0.04***	0.04***
	(0.004)	(0.003)	(0.002)	(0.002)
Log Distance (Other CBSA)	-1.4***	-1.4***	-2.0***	-2.1***
	(0.03)	(0.03)	(0.10)	(0.10)
Log Distance (Other CBSA) * Long-term CBSA	-0.10***	-0.09***	-0.16***	-0.14***
	(0.02)	(0.02)	(0.02)	(0.01)
Log Distance (Same CBSA) * Young Share		0.33***	0.33***	0.32***
		(0.04)	(0.04)	(0.04)
Log Distance (Same CBSA) * Low-income Share		0.15***	0.14***	0.13***
		(0.03)	(0.03)	(0.03)
Log Distance (Other CBSA) * Young Share		0.20***	0.13***	1.3***
		(0.03)	(0.04)	(0.12)
Log Distance (Other CBSA) * Low-income Share		0.04**	-0.05**	-0.44***
		(0.02)	(0.02)	(0.07)
Inclusive Value (Other CBSA)			-0.31***	-0.38***
			(0.05)	(0.05)
Inclusive Value (Same CBSA)			-0.59***	-0.44***
			(0.06)	(0.06)
Inclusive Value (Other CBSA) * Log Dest. CBSA Population				-0.04***
				(0.004)
Inclusive Value (Same CBSA) * Log Dest. CBSA Population				0.10***
				(0.03)
Inclusive Value (Other CBSA) * Young Share				0.46***
				(0.06)
Inclusive Value (Same CBSA) * Young Share				-0.12
				(0.22)
Inclusive Value (Other CBSA) * Low-income				-0.14***
				(0.03)
Inclusive Value (Same CBSA) * Low-income				-0.45**
				(0.19)
Fixed-effects				
Origin Time Tenure	Yes	Yes	Yes	Yes
Destination Time Tenure (same/other CBSA)	Yes	Yes	Yes	Yes
Fit statistics				
Observations	21,339,693	21,339,693	21,174,490	20,968,574
Pseudo R^2	0.77052	0.77214	0.71821	0.71814

Clustered (ZCTA Group) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Note: This table presents the results of the PPML regressions. The dependent variable is the number of migrants of each tenure level moving from the origin ZCTA to the destination ZCTA in each year. "Distance" measures the distance between the origin and destination. "(Same/Other CBSA)" indicates that the variable is interacted with a dummy equal to 1 if the move occurs within the same CBSA, and 0 otherwise. "Low Income Share" represents the share of renters with annual incomes below $35,000. "Young Share" refers to the share of renters aged 15 to 35 years. "Dest. CBSA Population" denotes the population of the destination CBSA.

Table A8: Second Step of Estimation

	Dependent variable:				
	log(Stayers/Movers)		log(Stayer/Movers)$_t + \beta \times log(Movers/TotalPop)_{t+1}$		
	(1)	(2)	(3)	(4)	(5)
Long-term Residence	0.020	0.020	0.074***	0.074***	0.078***
	(0.016)	(0.016)	(0.003)	(0.003)	(0.003)
Long-term CBSA	3.001***	3.001***	0.642***	0.640***	0.638***
	(0.015)	(0.015)	(0.008)	(0.008)	(0.008)
Log Distance to CBD					−0.017***
					(0.004)
Low-income					0.068
					(0.065)
Young Share					0.107*
					(0.063)
Log Total Population					−0.042**
					(0.018)
Log Housing Units					0.033*
					(0.020)
Log Median Gross Rent					−0.032
					(0.036)
Log Median Value					0.119**
					(0.050)
($V_t^{move} - \beta \times V_{t+1}^{move}$)			−0.397***	−0.395***	−0.394***
			(0.010)	(0.010)	(0.010)
($V_t^{move} - \beta \times V_{t+1}^{move}$)× Young Share				−0.153**	−0.085
				(0.063)	(0.059)
($V_t^{move} - \beta \times V_{t+1}^{move}$)× Low-income Share				−0.062	−0.016
				(0.052)	(0.050)
log Rent	0.219***	0.414*	−0.457**	−0.399**	−0.544***
	(0.037)	(0.250)	(0.196)	(0.201)	(0.211)
log Rent × Low-income Share				0.403***	
				(0.152)	
log Rent × Young Share				−0.087	
				(0.171)	
IV	No	Yes	Yes	Yes	Yes
ZCTA Group FE	Yes	Yes	Yes	Yes	No
CBSA-Year FE	Yes	Yes	Yes	Yes	Yes
Clustered SE	Yes	Yes	Yes	Yes	Yes
Observations	27,378	27,378	21,156	21,156	21,156

Note: *p<0.1; **p<0.05; ***p<0.01

Note:

Table A9: Fixed Moving Cost Estimates Using Aggregate Data

	Dependent variable:	
	log(Movers/Stayers) in the same Neighborhood	
	(1)	(2)
Long-term Residence		-2.616^{***}
		(0.016)
Low Income Share		-1.265^{***}
		(0.053)
Young Share		0.327^{***}
		(0.071)
Constant	-1.732^{***}	-0.863^{***}
	(0.011)	(0.010)
Clustered SE	Yes	Yes
Observations	24,252	24,252
Note:		*p<0.1; **p<0.05; ***p<0.01

Note: This table presents the results of regressions of the log ratio of movers within a neighborhood to stayers within the same neighborhood, with one observation per neighborhood-year. The structural equation is equation (18).