FEDERAL RESERVE BANK
PHILADELPHIA

Working Papers
RESEARCH DEPARTMENT

Fast Locations and Slowing Mobility

Patrick Coate
National Council on Compensation Insurance

Kyle Mangum
Federal Reserve Bank of Philadelphia

WP 26-26

PUBLISHED
May 2026

ISSN: 1962-5361

DOI: https://doi.org/10.21799/frbp.wp.2026.26

Fast Locations and Slowing Mobility

Patrick Coate* Kyle Mangum†

February 2025

Abstract

This paper shows the declining trend in internal migration in the United States is primarily due to increasing home attachment in "fast locations," areas with relatively high rates of population turnover. These locations were population growth destinations in the 20th century, with transient populations that settled as regional population growth converged. The qualitative patterns of the U.S. experience can be generated by a model of location choice in heterogeneous regions with overlapping generations when the population has a home bias that varies endogenously with the history of population change. Using a novel measure of home attachment, this paper estimates a structural model of migration that distinguishes moving frictions from home utility. Simulations quantify channels of the mobility decline. Rising home attachment accounts for much of the decline, predominantly in fast locations. Population aging explains most of the remainder but in a more spatially neutral way.

Keywords: declining internal migration, labor mobility, home attachment, rootedness, local ties, conditional choice probability estimation

JEL codes: J61, R23, R11, C50

*National Council on Compensation Insurance. Contact: pecoate@gmail.com.

†Corresponding author. Federal Reserve Bank of Philadelphia, 10 Independence Mall, Philadelphia, PA, 19106, USA. Contact: kyle.mangum@phil.frb.org.

1 Introduction

Internal migration rates in the U.S. have steadily trended downward in recent decades. The decline is pervasive across demographic strata, housing tenures, and household and family types (Molloy et al. (2011), Kaplan and Schulhofer-Wohl (2017)), suggesting a fundamental shift in the allocation of human capital across regions. This observation is alarming for policymakers because migration is considered a primary labor market adjustment mechanism. Americans have typically been regarded as a mobile population (Moretti (2012)), pioneers always in search of better opportunities, and there is rising concern that America has "lost its mojo" (Thompson (2016)). While a general notion of declining dynamism has entered the Zeitgeist of the academic and popular social sciences, the causes of the migration decline remain poorly understood, making it difficult to decide whether policy intervention is appropriate.

This paper shows that declining migration in the U.S. is driven by increasing rates of home attachment in the population. Social scientists have long found that households are far less likely to leave an area considered "home," where they have deep "roots" or social connections, and in the economics literature, home attachment is often prominent in studies of individual migration. This paper extends the notion from the individual to the population level by introducing the concept of *spatial demographics*: where people live and from where they came. When home attachments affect migration propensity, the amount of migration in an economy will be driven in part by the spatial demographics of its population, which in turn are a product of its history.

The analysis begins with the introduction of an important but overlooked fact: The recent decline in migration is largely due to declines in mobility out of "fast locations," places with higher-than-average rates of population turnover.[1] That is, migration has declined nationally because of convergence of fast locations' rates to those of slower locations. The paper then documents a set of stylized facts showing how the geographic history of population change would lead to first diverging and then converging levels of migration across places. Fast locations, mainly in the West, were the major destinations for population growth in the U.S. in the 20th century. Shifts in population across the geography of the U.S. left a large fraction of people in a places where they had shallow roots. Consequently, they frequently changed locations for myriad idiosyncratic reasons, making the growing places areas of high overall population turnover. As the pace of population shifts abated in the last decades of the 20th century, more and more people developed deep roots, and turnover decreased.

We then develop a model to interpret these facts and empirically test the hypotheses. The framework is a discrete choice model of location selection, in which agents are incentivized by each location's quality, their individual preference for their initial location ("home"), and

[1] Our analysis focuses on metro areas, according to our definition of a local labor market. The patterns we show are also apparent at the state level.

moving costs and idiosyncratic preferences.

First, we use a stylized version of the model to conduct a thought experiment on regional transition. In the stylized model, a regional economy features overlapping generations of agents who are endowed at birth with a home location. The economy experiences shocks to location quality that cause a shift of population across space, mediated across generations. Along the transition path, the population reallocation lowers the average home attachment, leading to higher levels of gross migration. As population growth converges, rootedness increases, and migration rates decline. These dynamics are in line with the U.S. experience. Notably, a model with home attachments – and especially home attachments that vary in intensity based on a location's history of population change – shows a greater level of gross migration and slower convergence than a model with moving costs but no home attachments.

The last part of the paper conducts a more rigorous measurement of the contribution of home attachment to the change in U.S. domestic migration. The empirical model extends the stylized model to account for heterogeneity in individuals (e.g., age and education) and the forward-looking nature of a costly migration decision. Choices are nested between moving or not, and if moving, where to migrate. The model is estimated on recent data depicting migration behavior by location, birthplace, and demographic group, aiming to separate home preferences from generic moving costs. A key part of identification of home preference is the extent to which people away from home return back to it – not only whether people at home fail to migrate – which indicates a preference for home over alternatives. For measurement of the intensity of home attachment, we introduce a statistic based on the probability that a birth cohort in a location had locally based parents. This encapsulates information on the regional history of population change and helps the model fit a prominent geographic feature of the data: that fast locations have more transient natives. We show how to tractably estimate a dynamic discrete choice model using a cross-section of choice probabilities, despite the nested structure of the model, an approach that may find use in other applications.

The results indicate a consistently large effect of home on the out-migration propensity and in-migration destination selection across a swath of demographic types. As a rough benchmark, being at home decreases migration propensity as if home were thousands of miles closer than other alternatives. The baseline model simulation can generate the spatial heterogeneity in in locations' migration rates, the age and education profile of migration, the propensity of natives to migrate less than others, and the spatial heterogeneity in migration among those at home.

We then simulate the model to see if it can produce a decline in migration in line with the empirical facts. The model is simulated backwards, using decades-old data to project migration as a function of individual and spatial demographics in years past. The simulation is contrived for the sake of hypothesis testing: We deliberately assume move costs and location quality do not change in order to see how far personal and spatial demographics can go in generating a mobility decline.[2]

We find a predicted trend is qualitatively and quantitatively in line with the observed decline and fits the spatial pattern. Decomposition exercises show that population demographics – in particular, the aging of the population – have contributed to the total migration decline, but spatial demographics are necessary to fit the size, spatial pattern, and ubiquity of the decline. Moreover, home attachment and aging interact to generate a larger decline in migration than the sum of the two trends.

The finding that home attachment is central to the migration decline is significant for several reasons.[3] First, though not the only factor in play, attachment has effects of a sufficient magnitude to warrant more attention from researchers. Second, it fits the otherwise peculiar spatial pattern of the migration decline. Third, it is pervasive across many different types of households, a puzzling feature for the literature to date.

The normative implications of our findings are mixed. On one hand, the welfare implications of the decline in mobility are actually quite positive, and a high migration rate may not be a desirable end in itself. More people are finding it optimal to stay in place and enjoy the benefits of social connections. Regionally, the migration slowdown is occurring in growing areas, not distressed, "left behind" locations. Moreover, the reasons for the decline are fairly mundane: The U.S. went through a spatial population transition, a particular phase of history, and settling into long-run equilibrium simply takes time.

On the other hand, an economy of higher home attachments is less nimble and less able to respond to shocks. Cyclical fluctuations may be more severe, mediated through labor market hardship rather than relocations. And, while the westward expansion of the U.S. may be over, there is some concern that population convergence is neither natural nor desirable. If population changes should be happening today but are not because of high costs of entry to productive and high amenity locations, then misallocation will result. These are interesting issues worthy of future study.

The paper proceeds as follows. We next review related literatures. Then we begin our empirical analysis in Section 2 with a thorough description of historical trends in migration and spatial population dynamics. Based on these findings, Section 3 writes down a basic model of migration amid regional transformation. Section 4 extends the basic model to a quantitative framework. Section 5 describes our estimation strategy, and Section 6 reports estimation results and compares model specifications. Section 7 reports on the baseline simulation. Section 8

[2]Psychological move costs are difficult to measure and usually specific to a model context. Assuming they increased could obviously generate a decline, but it also runs counter to intuition. We do conduct a simulation in which locations change quality over time in line with the model's estimates, but the effect on total gross migration is slight.

[3]To be clear, we do not claim the primitive preference for home has changed, but that the average exposure to home attachment has increased.

discusses simulations of the model over time and shows how various factors contribute to the mobility decline. Section 9 concludes. An appendix provides further details of the empirical methodologies and extensions of the model.

1.1 Related Literature

This paper provides a new explanation to a long-standing puzzle in the labor and spatial economics literatures, providing empirical support through new descriptive evidence and a structural model that rigorously quantifies the channel. Studies of the migration decline have looked for a structural mechanism in the economy, something pervasive across demographic strata. Home attachments fit this profile but have been overlooked because the spatial pattern of the decline has been ignored.

The national decline was noticed by researchers over a decade ago. An early herald was Fischer (2002), who noted that migration (i.e., moving one's local labor market) peaked around the 1970s and 1980s, while residential mobility (i.e., moving between homes within a labor market) had been steadily trending downwards for much longer. After Fischer, much of the work that emerged in the wake of the Great Recession (Molloy et al. (2011), Cooke (2011), Cooke (2013), Kaplan and Schulhofer-Wohl (2017)) emphasized the secular nature of the decline, finding compositional and cyclical explanations insufficient in magnitude and scope, although somewhat important in their own rights.[4]

The secular trend presented a puzzle and a cause for alarm. As labor mobility is thought to be one of the primary shock-adjustment mechanisms for regions[5] and individuals,[6] a natural concern arose that low mobility will result in spatial mismatch and lower aggregate productivity.[7] Some papers asked whether population responsiveness to local shocks was waning (Partridge et al. (2012), Dao et al. (2017)).

The pervasiveness of the trend led many researchers to posit structural changes in the economy. Hypotheses included the rise in dual-earner households and growth of communication technology (Cooke (2013)), information availability and the comparability of occupational returns across markets (Kaplan and Schulhofer-Wohl (2017)), spatial dispersion in income and home prices (Bayoumi and Barkema (2019)), and an aging workforce (Karahan and Rhee (2014)). Our analysis does not rule out these explanations, but none of them contends with the decline's

[4]To our knowledge, the only study to parse the migration changes across space is Frey (2009), although Frey focused on the cyclical dynamics of net migration in the early 2000s instead of the long run secular trend in gross migration in focus in our analysis.

[5]See, e.g., Blanchard and Katz (1992), Bound and Holzer (2000), Carrington (1996), Zabel (2012), and Cadena and Kovak (2016).

[6]See Topel (1986) or Kennan and Walker (2011).

[7]It is also worth noting that a substantial literature is devoted to understanding why labor mobility is slow or stagnant and not always in the expected direction (see, e.g., Sjaastad (1962), Lkhagvasuren (2012), Notowidigdo (2011), Autor et al. (2013), Dao et al. (2017)), and Yagan (2019)).

geographic pattern, which is the center of our paper.[8]

Our paper adds to the economics literature on nonpecuniary migration incentives. The state of the art model features idiosyncratic preferences, preference for initial or past locations, and move costs.[9] In particular, there is growing evidence for the role of social capital,[10] which also has a substantial literature in the population sciences outside economics.[11] Zerecero (2021) also estimates a home premium in a model with move costs, using French data and leveraging wage information in addition to moving probabilities. Zerecero (2021) finds reluctance to move away from home, a proclivity to return to it when away, and a wage discount for remaining home; all of these findings are consistent with home attachment.

Our findings show how these spatial demographics matter for aggregate population dynamics. Thus, home attachment, an important feature of many of the microeconometric studies of migration noted above, can accumulate to affect regional adjustment processes. Other research including move costs and/or home attachment is typically interested in their effects as frictions to labor supply (e.g., Rappaport (2004), Diamond (2016), Piyapromdee (2021), Heise and Porzio (2019), and Zabek (2024)). Our study is the first to connect the idea to the migration decline in the U.S. Our paper is distinct from Zabek (2024), which also incorporates a notion of home attachment, in that our objective is to explain the trend decline in migration, not the responsiveness of the population to labor market shocks, and in doing so we estimate the size of the attachments.[12]

More broadly, both Zabek (2024) and our paper highlight the nuanced intertemporal connection between gross and net population flows. While it is intuitive that gross population flows are related to net changes, we also show that the history of net change can affect the level of current gross flows. Therefore, our results are also related to the literature on regional convergence (e.g., Glaeser et al. (2006), Saks (2008), Gyourko et al. (2013), Ganong and Shoag (2017), Herkenhoff et al. (2018), and Hsieh and Moretti (2019)).

To focus on the quantification of direct, first-order effects of home attachment on migration trends, we employ a structurally estimated partial equilibrium model. The partial equilibrium

[8]By similar reasoning, our paper, which offers a geographic explanation to a geographic phenomenon, is only loosely related to the literatures on other forms of labor market dynamism that are often associated with the migration decline, but are not necessarily geographic in nature (Molloy et al. (2017), Molloy et al. (2016), Davis et al. (2012), Decker et al. (2016), Hyatt and Spletzer (2013)).

[9]See Kennan and Walker (2011), Bayer et al. (2009), Moretti (2011), Coen-Pirani (2010), Lkhagvasuren (2012), or Diamond (2016).

[10]See Dahl and Sorenson (2010), Munshi (2003), Glaeser et al. (2002), Kan (2007), David et al. (2010), Falck et al. (2012), Alesina et al. (2015), Büchel et al. (2020) and Koşar et al. (2021).

[11]See, for example, Dawkins (2006), Michielin et al. (2008), Mulder and Malmberg (2014), Belot and Ermisch (2009), Clark and Lisowski (2019).

[12]That the intensity of home attachment can vary with past population dynamics is also novel to our study. Zabek (2024) is concerned with the within-population heterogeneity in the strength of attachment, such as the selection effect that arises when only those with the strongest preferences remain in economically depressed areas, but this is a selection effect, not due to path dependence.

setting eases some demands on the model – there is no need to assume the data are in steady state, for instance, or that workers are myopic. In our structural estimation, we are taking the regional evolution as given, using a cross-section of data and spatial variation for identification, and measuring how the past evolution affects migration rates into today.

On the technical side, our paper contributes to the literature on the estimation of dynamic discrete choice models. Since Kennan and Walker (2011), recent work has made progress in making estimation tractable even with large state and/or type spaces. A close comparison to our study is Davis et al. (2021), which utilizes odds ratios in a conditional choice probability (CCP) estimation[13] of neighborhood choice among a single metro area.[14]

The distinct feature in our case is the nesting of the decision model, a formulation inspired by Monras (2018), which complicates the formulas for CCPs. We can show that with suitable adjustments to the conditional logit model, concise estimating equations can be obtained. While our solution was custom-built to solve our problem, we hope that our approach can open up new applications of estimable nested dynamic discrete choice models.[15]

2 The Migration Decline: Its Geography and History

We first introduce a set of novel empirical facts that help motivate, formulate, and preview the results of the model.

2.1 Data Overview and Preliminaries

This study relies on an assemblage of data from several different publicly available sources. We briefly describe them here and provide more details in Appendix A.

[13]For seminal work on CCPs, see Hotz and Miller (1993) and Arcidiacono and Miller (2011). For recent applications, see Bishop (2008) and Ma (2019). Though not explicitly characterized as CCP estimation, Artuc et al. (2010) derive an estimating equation from a choice value function.

[14]Bayer et al. (2016) also estimate a neighborhood choice model with adjustment costs, although the approach to dynamics is to derive formulas for net present values instead of utilizing CCPs. The aforementioned Zerecero (2021) also writes down a dynamic model and uses the exact hat algebra approach of Caliendo et al. (2019) to account for continuation values in the migration decision. Zerecero (2021) also addresses identification when flow data are sparse.

[15]Nesting is a flexible and generalizable procedure. A few examples with similar structure to ours, with a binary and then multinomial choice, include the following: whether to enter the labor market, and which job to work; whether to attend university, and which field to study; whether to search for a new job; and in which markets to do so; whether to change occupations or industries, and which sector to enter; whether to export, and which country or product to sell. We also note that nested decision problems have been employed in quantitative spatial models, such as Couture et al. (2024), Tsivanidis (2023), and Baum-Snow and Han (2024).

2.1.1 Data Sources and Uses

Our measures of migration come from two sources. The first is the migration flows tables for 1991 to 2019 from the U.S. Treasury's Internal Revenue Service (IRS). The IRS infers migration events from changes in the address on individual tax returns for two successive years, publishing the total county-to-county flows in each year, as well as the total stayers in, inflows to, and outflows from individual counties. Because of its consistent reporting (annual since 1990), geographic detail (counties), and its large sample size (a near universe of taxpayers), the IRS data are the principal source for describing trends in mobility across different geographies.[16] However, the data lack demographic information that is essential to a microfounded model of migration flows.

The second source for measuring migration flows is the American Community Survey (ACS) for 2005 to 2019, obtained from Integrated Public Use Microdata Series (Ruggles et al. (2019)). The ACS microdata contain rich demographic information such as age, education, and – importantly for our purposes – place of birth, so it will be the primary source for quantifying our model. The ACS reports the respondent's current and one-year-ago Public Use Microdata Area (PUMA) of residence, from which we can elicit migration probability (move or not) and direction (origin-destination).

The 1990 and 2000 censuses also report a retrospective migration question, but at a five-year instead of one-year look-back, which complicates comparison because return and repeat ("onward") moves are frequent (Kennan and Walker (2011), DaVanzo (1983)). We derive one-to-five-year conversions using location histories in the longitudinal Panel Study of Income Dynamics (PSID) (Institute for Social Research (2021)) to measure return and onward migration rates by age, education, and birthplace status (see Appendix A.4). This provides corroborating evidence for our stylized facts, but because it is based on assumptions about the conversion procedure, we do not rely on the five-year rates for the model.

Finally, for historical measures of population, we use harmonized decennial census records from the National Historical Geographic Information System (NHGIS) (Manson et al. (2018)), and for its distribution by age, education, and birthplace, we use microdata samples from 1880 to 2000 from Ruggles et al. (2019).

2.1.2 Geographic Concept for Defining Migration

An initial issue is defining at what distance a move becomes "migration." We use a geographic concept similar to a Core-Based Statistical Area (CBSA) or commuting zone (CZ) that

[16] The IRS data underwent a change in method in the 2011-2012 tax year that resulted in noticeable, but as-yet unexplained, differences in the sample represented and possibly the migration concept, sowing doubt about the reliability of the data for these later years (DeWaard et al. (2022)). We present the data for the period 2012-2019 but rely only on the consistent sample of 1990-2011.

Figure 1: Local Labor Market Migration Rates, 1991-2017

One-Year Rates Over Time **In and Out Rates (IRS)**

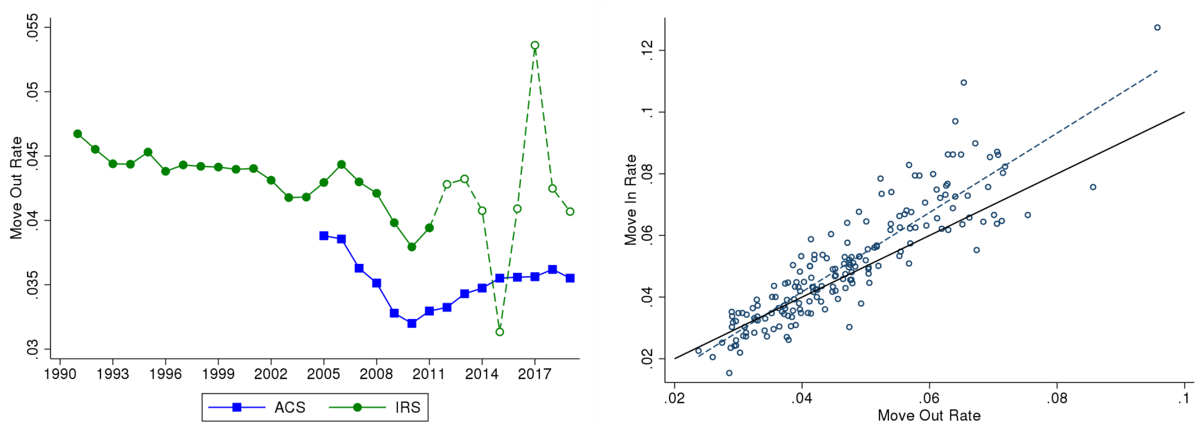

NOTES: The lefthand figure plots the one-year out-migration rate over time as defined by LLM and observed in the IRS and ACS microdata samples. The data become dashed series in 2012 and following to reflect a break in the method the IRS used to define move rates. The righthand figure plots average LLM in- versus out-mobility rates in the early IRS samples. (Source: IRS and ACS data.)

we term a local labor market (LLM). Intuitively, this functions like a CBSA or CZ, but allows us to standardize the geographic boundaries over time, important for working with historical data. We can map PUMAs and counties into LLMs for each year of data, 1880 to 2019. Migration is defined as exiting one LLM for another, and moves within the boundaries of an LLM are non-migration events. We focus on the urbanized LLMs.

We prefer the LLM definition to state because LLMs are economic, not political, areas.[17] However, when defining one's home location, we are confined to use the state definition because the ACS and censuses ask for respondent's state of birth. We account for multi-state LLMs and have examined sensitivity of results to the alternative mappings of state to LLM in the assignment of home.

Using the LLM concept, Figure 1 lays out two initial matters. The lefthand plot displays the national average migration rate over time in the IRS and ACS data samples. There is a clear downward trend (amid some cyclical fluctuations) best seen in the longer IRS time series. The drop from 1991 to the 2005-2011 average was 0.5 percentage points, or a 20 percent drop in the rate at which households change LLMs. Following Molloy et al. (2011) and Kaplan and Schulhofer-Wohl (2017), our emphasis is on the downward trend and not the cyclical fluctuations. Interestingly, the trend seems to taper in the 2010s (a feature which our model simulations will replicate).

[17]This is one major reason we do not rely on the Current Population Survey (CPS), which provides only state of residence. Another is that the CPS sample is too small to split into subgeographies crucial to our analysis. The CPS is the outlier dataset, registering a migration decline much larger than other sources, and exhibits inconsistencies calling into question its reliability for measuring migration rates over time Hyatt et al. (2018).

Next, we explain the language of "fast" vis-a-vis "slow" locations. The righthand plot scatters out-mobility to in-mobility rates for LLMs in the 1991 IRS data. Gross migration rates are correlated – locations with high degrees of inflow also exhibit high degrees of outflow – and there is a large variance across places in the degree of turnover.[18] Our use of fast, medium, and slow categories is based on population turnover at the start of the data in 1991. Each category has one-third of U.S. population to the closest approximation. From here, we will be grouping LLMs by category, but Appendix B contains scatterplots of many of the following analyses.

2.2 Fast Locations Drive the Migration Decline

The first novel fact is that the decline in migration is heterogeneous across space. Specifically, high-turnover fast locations are primarily responsible for the national decline by converging towards their slow counterparts. Figure 2 exhibits the annual out-migration rates for LLMs split into terciles by their mobility rate – fast, medium, and slow. The lefthand panel shows the one-year rate from the IRS. The most mobile third of cities show a strong downward trend, dropping from about 5.7 percent to 4.6 percent from 1990 to 2011 (a 21 log point decline). The change for the middle third was much smaller, declining from about 4.3 to 4.0 percent (a 7 log point decline). The least mobile third showed almost no decline.[19]

The trend from fast locations is similar when deriving the migration time series from microdata. The righthand panel displays five-year rates using our one-to-five-year conversion procedure for the ACS-era years (2005-2019). That the proportional changes are so similar in census and IRS data provides strong corroborating evidence that the fast/slow distinction is not an artifact of a particular dataset. Notably, the three groups of LLMs have similar cyclical fluctuations but disparate trends. And in each, the trend appears to taper in the 2010s.

Understanding the decline in outflows from fast locations is then essential for studying the national mobility decline. A simple accounting exercise helps fix ideas. By definition, the fast tercile of locations makes up one-third of the population but a greater share of out-migrants – about 44 percent of them in the early 1990s. We project how many migrants there would be in 2010 if all cities moved at 1990 levels, and we then take the difference from actual migrants as the number of "lost migrants." This accounting indicates that, with a larger share of migrants and a larger decline in rates, the fastest third of LLMs is responsible for 64 percent of the national decline.[20] In contrast, slow locations account for only 13 percent of the decline despite

[18]This fact is not new (see seminal work by Ravenstein (1885) and Sjaastad (1962), and more recently, Coen-Pirani (2010)), but we are the first to tie it to trends in mobility. It is important to note the distinction between places concerns *turnover*. That is, high rates of out-mobility are not signs of a population exodus. Indeed, the fit line from in- to out-migration is slightly steeper than the diagonal, indicating that on average, higher turnover places have net in-migration.

[19]The differences in LLM category trends are statistically significant by standard measures and are apparent among many types of destinations.

[20]As a particularly notable example, the cities of California make up 31 percent of the lost migrants, and

Figure 2: Changes in Out-Migration Rate Over Time by Average Initial Turnover

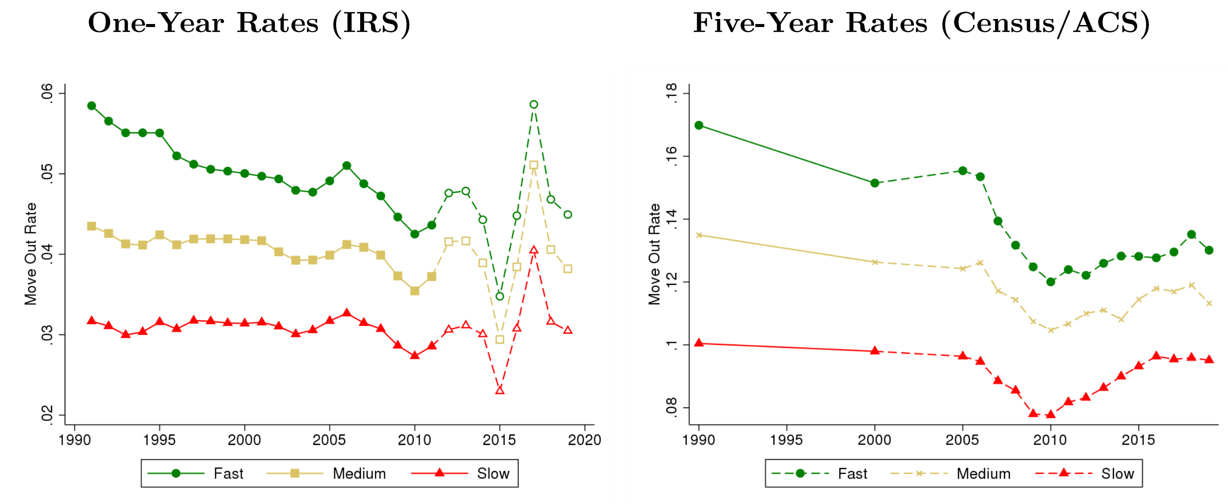

NOTES: The lefthand plot shows the annual out-migration rates in the IRS data by initial mobility tercile. The data become dashed series in 2012 and following to reflect a break in the method the IRS used to define move rates. Figure B1 displays the data in scatterplot. (Source: IRS data.) The righthand plot shows the five-year out-migration rates by initial mobility tercile in the census, spliced with the implied five-year migration rates using ACS migration data adjusted with the repeat migration model from the PSID. The data become dashed series after 2000 and following to reflect the switch from actual to implied five-year rates. The conversion is discussed in Appendix A.4. (Source: ACS, Census, and PSID data.)

comprising 40 percent of population.[21]

The spatial pattern invites two natural questions: What is different between fast and slow locations, and what is changing about them?

2.3 Local Labor Market Attributes

As a first pass, we look for correlations of attributes of LLMs to migration rates and migration decline, summarized in Table 1. We run one regression for each variable, reporting the coefficient, standard error, and R^2. We then combine into a multivariate regression. We look at basic attributes: population size and growth, population demographics of age and education, and moments of the wage distribution (mean and standard deviation).

Individually, these variables have little explanatory power for distinguishing the mobility rate of an LLM. The share aged under 40 and the share with a college degree have the expected sign with migration rates but have little explanatory power. Moreover, they are correlated with a greater decline in migration. These are clearly important predictors at the individual level but do not seem determinative at the LLM level.

Higher mean incomes appear negatively correlated to migration rates, but also with little explanatory power and mixed results on the decline. The dispersion of income is positively correlated with migration rates and also with larger declines, but most of this is weighted onto

Southern California alone - Los Angeles, Riverside/San Bernardino, and San Diego - makes up 18 percent.

[21]The single largest LLM, New York City, straddles the boundary of the bottom third of mobility, with the associated lumpiness making the bottom tercile actually more than one-third of the population.

Table 1: Correlation of Mobility Rates and Local Labor Market Attributes

	Dependent Variable: Initial Migration (1991)			Dependent Variable: Change in Migration (1991-2010)		
Regression:	Single		Multi	Single		Multi
LLM Attribute	Coef (se)	R^2	Coef (se)	Coef (se)	R^2	Coef (se)
ln Population Size	-0.333	0.060	-0.774	-0.029	0.002	0.219
	(0.097)		(0.113)	(0.050)		(0.059)
Population Growth	3.549	0.246	3.429	-1.264	0.127	-0.816
	(0.462)		(0.462)	(0.246)		(0.240)
Aged Under 40	8.171	0.023	6.446	-9.508	0.124	-7.159
	(3.991)		(3.418)	(1.876)		(1.779)
College Educated	0.817	0.001	6.998	-2.186	0.033	-3.950
	(1.788)		(1.890)	(0.873)		(0.984)
Mean Income (Noncollege)	-2.228	0.046	0.443	0.447	0.008	1.079
	(0.751)		(1.068)	(0.380)		(0.556)
Mean Income (College)	-0.894	0.005	1.080	-1.155	0.034	-1.826
	(0.929)		(1.258)	(0.454)		(0.655)
Income Dispersion (Noncollege)	8.498	0.079	10.374	-5.637	0.141	-6.089
	(2.160)		(1.964)	(1.035)		(1.022)
Income Dispersion (College)	7.561	0.038	4.604	-3.003	0.025	-1.713
	(2.818)		(2.335)	(1.409)		(1.215)
Combined R^2			0.486			0.435
LLMs	183	183	183	183	183	183

NOTES: The table reports regression coefficients of local labor market attributes to migration levels, by row, and changes, by column heading. The first column under each heading reports coefficients from single variable regressions, followed by the implied R^2. The last column in each subheading reports the multivariate regression. Declines in migration are negative changes, so a negative correlation indicates a larger value of the variable is associated with a larger magnitude decline. (Source: IRS migration data; ACS data for demographic and labor market attributes; Census county population data for population estimates.)

the noncollege educated group. While we find this somewhat interesting,[22] the local labor market attributes do not appear to be greatly important for explaining migration rates at the LLM level (though they may well be important for determining the direction of some flows). Hence, for the rest of the paper, we will de-emphasize these local labor market attributes and, in the model, generically account for "location quality."

The single strongest correlate is population growth. Its R^2 in a single variable regression is half the explanatory power of all the attributes combined. As we will see, this stems from the spatial demographic history of growth in these locations. To understand this more clearly, we first elaborate on the importance of home attachment in migration.

2.4 The Importance of Home in the Migration Decision

The clearest connection of the migration decline is to population growth. Why would growing places be fast, and why does their outmigration decline after they grow?

A critical fact to establish at the outset is that one's home location occupies a special status in the choice set; that is, home offers a utility premium not available elsewhere. There is a substantial literature in the social sciences documenting the importance of "home" (broadly defined) in determining migration decisions, so while this idea is not new, we show here that

[22]In unreported analysis, we incorporate a model extension whereby income dispersion affects gross migration rates. The net effects on migration decline were relatively small, so we leave this for future investigation.

the measures of home available in the ACS and census are predictive of migration propensity in the expected way.

Table 2 uses ACS data to report annual mobility rates by age and education in total and disaggregated by birthplace status. The available variables in the census give a coarse but consistently defined measure of "home." We define home status hereafter as (1) born in a state represented by their current LLM ("At Home"), (2) born in some other U.S. state ("Other U.S.," or simply, not at home), or (3) born outside the U.S ("Outside U.S.").

Some well-known patterns appear: The young are more mobile than the old, and the college educated are more mobile than noncollege, especially in youth. But among all categories, there are major differences by birthplace status: Those living away from their birthplace are several times more likely to migrate than those at home. The foreign born's internal migration resemble those living at home, although among the college educated, mobility rates for the foreign born are somewhat closer to the away-from-home rate.

It is important for interpretation to understand whether the difference in move rates by at-home status is due to an actual utility-enhancing component – a *preference* for home – producing strong attachment to the place, or if the gap between columns 2 and 3 is due to selection on move costs among those who have never left their initial places.

One testable hypothesis is that if home is deemed especially valuable as a location attribute, then home will be chosen more frequently as a destination when living away from it. Column 5 reports the conditional choice probability of moving home when living away from it and when migrating somewhere. Roughly one-tenth of moves are returns home.[23] For comparison, in column 6 we report the average probability in the general population of migrants of selecting one's home – the expected probability if home were not special.[24] These are all orders of magnitude smaller; a destination is 9 to 21 times more likely to be chosen when it is one's home.[25]

These patterns provide initial evidence of the existence of home attachment via a utility premium, and we will provide additional evidence after developing further concepts of home attachment.[26] This explanation is important to bear in mind as we review the spatial pattern

[23]The precise measurement is subject to definition of home LLM when there are multiple LLMs in one's birth state. In Table 2, we use a conservative definition of migration to home, using weights to predict which LLM in the birth state could be home, and assuming moves to other LLMs are not moves home; for example, Los Angeles and San Francisco cannot both be defined as home in our calculation. Including any home-state move roughly doubles the estimate presented here.

[24]A random choice probability, one out of 275 LLMs, is about 0.36 percent. In column 6, we are adjusting the probability for the relative sizes of the LLMs. For example, because of New York's market size, there are many New Yorkers living about the country who may return home, and it is a popular destination for migrants from all birthplaces.

[25]Zerecero (2021) finds a similar return-to-home pattern in French data.

[26]Appendix C provides additional results on the effect of home on migration propensity utilizing the ACS microdata.

Table 2: Move Rates by Age, Education, and At-Home Status

	Move Out Rate (%)				Destination Choice Probability (%)	
Birth Status: Education-Age	Total 1	At Home 2	Other U.S. 3	Outside U.S. 4	Actual Home 5	Synthetic Home 6
Noncollege						
20s	5.79	4.41	12.01	4.12	14.45	0.70
30s	3.52	2.62	6.89	2.66	11.86	0.68
40s	2.35	1.68	4.40	1.77	10.89	0.64
50s	1.91	1.29	3.41	1.53	9.35	0.57
College						
20s	10.02	7.33	15.02	9.89	12.47	1.10
30s	5.07	3.09	7.57	5.68	10.06	1.00
40s	2.56	1.52	3.77	2.70	7.79	0.91
50s	2.20	1.41	3.16	2.03	6.96	0.73

NOTES: The table reports mobility rates by birth place. A respondent is in birthplace if residing in an LLM within his/her reported state of birth. Move home rates (column 5) are LLM move-in rates, conditional on a migration occurring, weighted by the probability the LLM is the respondent's birth LLM. The synthetic move home probability is a weighted average of conditional choice probabilities of moving into the home LLM for respondents not born in the LLM; for example, the probability of choosing Los Angeles by people not born in California. (Source: ACS data.)

of the mobility decline and the evolution of population that preceded it.

2.5 Fast Locations: Home of the Transient

Table 2 makes clear that nativity is a principal feature for predicting an individual's migration propensity. We next show a principal feature of fast locations is their loosely-attached populations.

Figure 3 shows population shares by home status within age and education categories, as in Table 2, for fast, medium, and slow LLMs. Within each age/education category, fast locations derive a larger fraction of their population from those in the not-at-home category and fewer from the at-home. This is seen best in the at-home category where a "stair-step" pattern clearly shows increasing fractions of natives as the LLM turnover rate declines.[27] Thus, among fast locations, a greater share of their population has arrived there from elsewhere, and consequently, a larger fraction of their population is more susceptible to moving (again).[28]

2.6 The Intensity of Home Attachment: Deep and Shallow Roots

There is another layer to home attachment, moreover. Fast locations not only have smaller shares of the at-home, but their at-home residents appear to be less attached to their native places. Figure 4 plots the migration propensity by birthplace status for each age and education

[27] Zabek (2024) includes maps of the U.S. depicting population "born locally."

[28] The intuition here is similar to Coen-Pirani (2010), which develops a model of migration over local labor markets in which new arrivals are more subject to mismatch shocks that result in out-migration yet again. That idea appears in the job search literature (Jovanovic (1979)). Our model is explicit in modeling the native preference attachment for initial places (including predisposition to returning home).

Figure 3: LLM Population Share by Birthplace Status, Across Age and Education

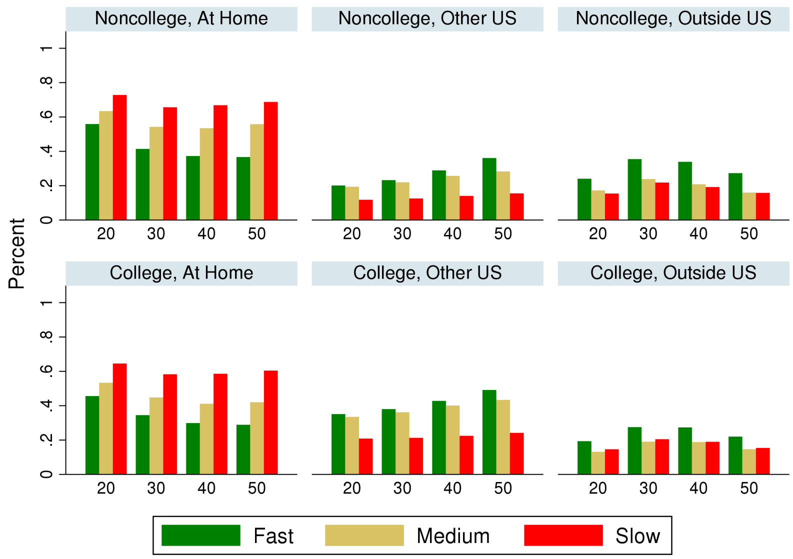

NOTES: The figure plots population shares by age and educational group and at-home status, split by LLM speed category. The shares are calculated across birthplace status as a percentage of the age/education group population in the LLM group (at-home, other U.S., and outside U.S. will add to 100 percent). Figure B2 presents the data in scatterplot. (Source: ACS data; IRS data for LLM categorization.)

group, as in Table 2, but split by origin LLM mobility tercile. Among those U.S.-born but not at home (middle column), fast, medium, and slow cities send away residents at similar rates, and patterns among the foreign born are mixed. However, among the at-home group, another stair-step pattern emerges: Natives of fast locations move away at higher rates than natives of slower locations.

Thus, the fast locations not only have more nonnative residents, but also send away their natives at higher rates. This suggests the *intensity* of home attachment varies across space. Why might natives be less attached in one place versus another? The literature reviewed in Section 1.1 indicates home is preferable in large part because of social and family networks, and hence stronger connections could produce stronger preferences. We seek a method for measuring the intensity of connections across natives in different places in a way that can be consistently applied over time.

The consistent reporting of respondents' state of birth in the census provides an avenue for constructing such a measure, which we will call "roots."[29] Appendix A.3 contains a detailed discussion of implementation. In brief, we measure the probability of being born to parents native to one's own location of birth, and then apply this measure to current generations by matching by birth year and place. For example, a 30-something in 2010 was born in the 1970s and was therefore under 10 in the 1980 census. We take children under 10 in 1980 living in, for instance, the Boston LLM, and then, using the family relationship variables in the census,

[29]We certainly did not invent the terms "roots" or "rootedness," as these are used in a variety of contexts across the population sciences, but we mean to use them here in a particular way.

Figure 4: LLM Migration Rates by Age, Education, Birthplace Status

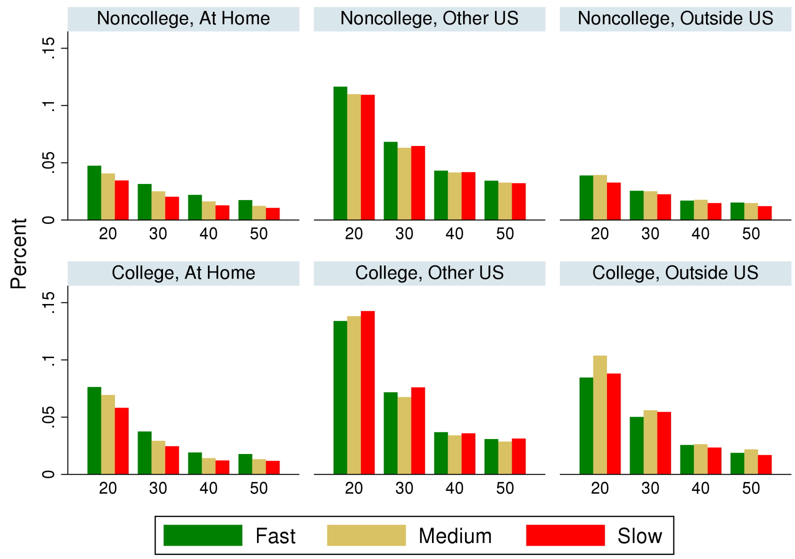

NOTES: The figure plots migration propensities by age and educational group and at-home status, split by LLM speed category. The national averages by person type are reported in Table 2. Figure B3 presents the data in scatterplot. (Source: ACS data; IRS data for LLM categorization.)

summarize the proportion of their parents who report a home state of Massachusetts. This proportion forms an index of a cohort's home attachment, e.g., how highly attached to Boston someone born there 30 years earlier is, even if living in another location by 2010.

Rootedness by our definition is feasible to measure and can be constructed consistently over time, but of course is only a proxy. However, depth of roots lines up well with the observed pattern of migration in Figure 4. In fast locations, average rootedness is 0.46, meaning that the typical at-home resident has just less than one parent also native to the place. Natives of slow locations, on the other hand, have average roots of 0.72, and medium locations register in between at about 0.69. These differences are consistent with natives' apparently lower degree of attachment in fast locations. In Appendix C, we provide additional evidence that rootedness is a good measure of the intensity of home attachment (and not spuriously correlated with other factors affecting mobility).

In summary, fast locations have high transience through two channels: higher shares of population born elsewhere and natives who are relatively less attached to their homes. How and why did fast locations become the home of such transient populations? We next review the history of population growth that would cause the current pattern to arise, with implications for why migration has more recently declined.

2.7 Fast Locations: Centers of Continental Population Expansion

There is a clear regional component to LLM-level population turnover. For instance, there are no fast locations in the Midwest and only two in the Northeast.[30] In contrast, 87 percent of the West's population lives in a fast location.

The regional pattern hints at what is different about fast places: high population growth as the result of geographic expansion. Figure 5 plots decadal rates of population growth by region (in the lefthand plot, all places within region, regardless of metro status) and by LLM category (in the righthand plot, all metro areas, regardless of region). The 20th century was a time of continental expansion, with the West the engine of population growth, as sparsely populated arid places transformed to urban centers.[31] Cities such as Los Angeles, Phoenix, and Las Vegas burgeoned from small outpost towns with just a few thousand residents at the start of the 20th century to major urban areas at its end. The South emerged as a population destination in the latter half of the century, and the Frontier region (central plains, Texas, and Rocky Mountain states) retains a relatively high rate of growth. The Midwest and Northeast lagged throughout.

Expansion to these new regions was relatively sudden (by historical standards), as new technologies made developable areas that were once too remote or difficult to inhabit at a large scale. The reasons for the growth of these areas are varied, but there are a few themes. One is development of water management technology (Luckingham (1984), Reisner (1993)). For example, the development of Los Angeles followed the completion of the aqueduct in 1913 and its increasing use for urban water delivery in the 1920s. On the opposite coast, the technology of water control aided Florida's development, as storm water was captured and swamps were drained (Barnett (2008)). Moreover, advances in transport technology and climate control made accessible newly desirable parts of the American continent. Railroads connected the population centers of the east to the west and the Florida peninsula (Wiggins (1995)). Later in the 20th century, air conditioning played an important role (Trippett (1979)). Besides making hot summers tolerable, air conditioning enabled the construction of high-density residential structures and large-scale industrial production. These technological trends meant the 20th century was a phase of history characterized by an opening of the American continent to urbanization at a scale not previously experienced – and one that had converged by century's end.

2.8 The Evolution of At-Home Status Across Space

The history of population growth had implications for the spatial demographics of the country. Figure 6 uses historical census microdata on place of birth and place of residence to report

[30] These are Washington, D.C., and Manchester, New Hampshire. New York has one of the slowest turnover rates, even adjusting for its size (larger cities tend to have lower turnover).

[31] See Chinitz (1986) for a discussion of American regional transformation.

Figure 5: Population Accumulation Over Time, by Region and by LLM Mobility Rates

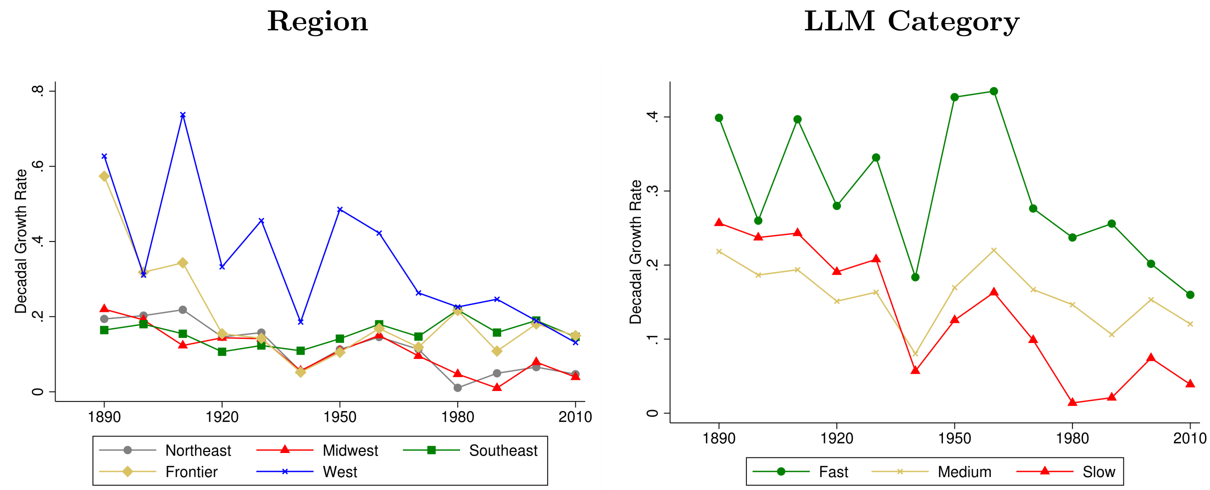

NOTES: The figure reports decadal population growth rates by LLM speed category (left) and region (right). The Frontier comprises the states of the central plains (including Texas) and Rocky Mountains. (Source: Census county population estimates, harmonized over time by NHGIS, and IRS data for LLM categorization.)

the proportion of population at home, not at home, or born outside the U.S. Separately, a dotted line reports the ratio of at-home population to other U.S.-born (dropping the foreign born from denominator) to measure the share of U.S. natives who are in their birth LLM.

Trends in immigration at the beginning and end of the century were similar nationally, but mid-century, the U.S.-born population reallocation tilted toward the West and Wouth, as seen in the upper right of Figure 6. This growth slowed, and the population share of the at-home gradually increased, from about 60 percent in 1960 to 80 percent by 2010. As one example, in 1960, 20 percent of the U.S.-born Los Angeles residents were from California; by 2010, that share was 71 percent. Medium and slow LLMs, in contrast, consistently had at-home shares at higher levels, on flat or downward trends.

The dynamics of home status plotted in Figure 6, resulting from the population boom and its waning as shown in Figure 5, offer an explanation for the transience of fast locations discussed in Sections 2.5 and 4.3. First, as the destinations of population growth in the 20th century, fast locations are home to more migrants from other birthplaces, meaning they have more of the "mover class" of not-at-home, U.S.-born residents. Second, because of the decades-long timing of the population transition, greater fractions of fast-location natives are born to not-at-home parents. Finally, the convergence in population growth rates in more recent decades has meant that both of these margins – home status and depth of roots – began to increase, making fast locations less transient in recent decades.

Figure 6: Population Share from Birthplace Source, by LLM Mobility Category

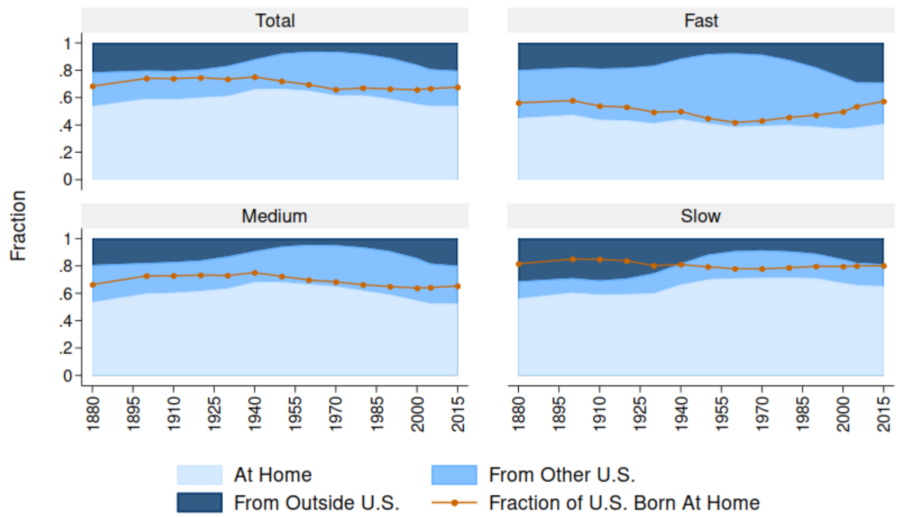

NOTES: The dotted line measures the fraction of U.S.-born population at home. At Home refers to living in an LLM in one's state of birth. From other U.S. refers to a birthplace in another U.S. state not covered by one's LLM. From outside U.S. are foreign born from any other country (or non-continental U.S. states and territories). The figure panels are summarized by LLM mobility categorization (fast, medium, and slow). Total includes all LLMs and rural areas/excluded small cities. (Source: Authors' calculations using census data; IRS data for LLM categorization.)

3 Regional Dynamics with Home Attachments

We next write down a model of location choice. We begin in this section with a stylized version of a model of regional dynamics with population cohorts. The purpose of the stylized model is to show qualitatively how shocks to locations, mediated across generations, affect the migration rate. In the subsequent section, we use an empirical model to test whether the home attachment mechanism has quantitative explanatory power in U.S. data.

3.1 Location Choice Model

An agent's objective takes the form of a discrete choice model,

$$V(o) = max_j[\tilde{v}_j(o)], \tag{1}$$

where j indexes possible locations and o indicates origin location. As is common practice, the choice specific values take a form separable in the location attributes and an idiosyncratic, iid shock to preferences,

$$\tilde{v}_{ij}(o) = v_j(o) + \varepsilon_{ij}, \tag{2}$$

where $v_j(o)$ are the value components of j common to all individuals with state o and ε_{ij} represent the iid preference shocks to individual agents i. When the shocks are distributed as Type I extreme value (T1EV) with scale parameter λ, there are convenient expressions for choice probabilities and continuation values. The choice probability is defined as

$$Pr(j|o) = \sigma_{j,o} = \frac{exp(v_j(o))^{\frac{1}{\lambda}}}{\sum_k exp(v_k(o)^{\frac{1}{\lambda}}}, \tag{3}$$

where λ governs the dispersion of the T1EV shocks.

The choice-specific common values are further separated into three components: (1) a location quality common to all agents, (2) a persistent location preference specific to a certain type of agent, and (3) moving costs representing frictions to changing locations. In the stylized version of the model, the first will be subject to shocks to represent "regional evolutions," the second will account for home attachment, endowed to agents based on their birthplace, and the third will be a standard moving friction:

$$v_{jo} = \mu_j + u_{jh} + c_{jo}. \tag{4}$$

The u_{jh} term is the home premium obtained in location j afforded to an agent born in h. Here (and later in the empirical model) the premium is present when living at home and absent everywhere else; i.e., $u_{jh} > 0$ if $j = h$, $u_{jh} = 0$ otherwise.

The last core feature of this model is a nesting of the moving decision. This feature is important for matching the data, and although it is not essential for the qualitative features of this stylized model, we introduce it here for completeness.

3.2 Nesting the Model

The reason for nesting the location choice model is to allow for differential elasticities of inflows and outflows to location attributes; see Monras (2018) for a discussion of the matter.[32] We need a nested decision model in our case in order to fit the large discrepancy in the impact of home status on migration move-out and conditional move-in propensities. Appendix C details the empirical case for this employing this feature in the current context. In short, the odds ratio for move-out probability between an at-home and not-at-home individual is much smaller in magnitude than the move-in probability to home versus other locations. On the move-out side, the odds ratio is approximately 0.1, while the odds ratio for choosing to return home relative to choosing somewhere else is orders of magnitude larger, roughly 5-10 depending on the group and measurement of home. A single parameter is not sufficient to meet this pattern.

In the nested version, two substages may occur. First, workers decide whether to stay in their current location. Second, and contingent on deciding to migrate, they choose a new location. The upper nest is the binary stay (s) or move (m) decision, and the lower nest is a choice among

[32]Zabek (2024) also notes that the share of people staying in their birthplace is only weakly related to the change in nonnative arrivals, indicating different elasticities even in the long run.

alternatives besides the origin.

Starting in the latter stage, assuming that a worker has decided to migrate, the choice-specific value of a destination is (4). The choice probability is like (3), but excluding the origin as an option:

$$Pr(j|o,m) = \sigma_{j,o|m} = \frac{exp(v_j(o))^{\frac{1}{\lambda}}}{\sum_{i\neq o} exp(v_i(o)^{\frac{1}{\lambda}}}. \tag{5}$$

The former stage has two values to compare: moving and staying. The value of moving is that of choosing a destination optimally in the latter stage – the expected value of $max_j\{\tilde{v}_{j\neq o}(o)\}$, which follows (2), but excludes the origin from the choice set:

$$V^m(o) = \lambda ln\big(\sum_{i\neq o} exp(v_i(o))^{\frac{1}{\lambda}}\big). \tag{6}$$

The value of staying (s) is simply maintaining the value of the current location,

$$V^s(o) = v_j. \tag{7}$$

With a T1EV draw on the upper nest decision, the respective probabilities are

$$Pr(stay|o) = \sigma^s = \frac{exp[V_s(o)]^{\frac{1}{\delta}}}{exp[V_s(o)]^{\frac{1}{\delta}} + exp[V_m(o)]^{\frac{1}{\delta}}} \tag{8a}$$

$$Pr(move|o) = \sigma^m = \frac{exp[V_m(o)]^{\frac{1}{\delta}}}{exp[V_s(o)]^{\frac{1}{\delta}} + exp[V_m(o)]^{\frac{1}{\delta}}}, \tag{8b}$$

where the elasticity of the upper nest is governed by parameter δ. (We use superscripts to denote s, m in the stage 1 decision and subscripts to denote locations in the stage 2 decision.)

The probability of choosing a location is then: (i) for the current location, the probability of staying, σ^s or (ii) for a new location, the joint probability; that is, the product of the probability of moving and the probability of a mover choosing the new location, $\sigma^m\sigma_{jo}$.

3.3 Regional Shocks

We can use this version of the model to illustrate the potential effect of home attachment on an economy undergoing regional change. We run the following thought experiment: what is the path of an economy's migration rates if some regions experienced an increase in attractiveness? How would the migration path differ if there was a preference for home?

The economy is a closed set of J locations, with an individual location indexed j. Time is discrete, with periods indexed by t. The population is comprised of A cohorts of households. Cohorts are born at age $a = 1$ and progressively age one by one, $a_{t+1} = a_t + 1$, each period t until their certain death at age A. At birth, each agent is endowed with a home location, which is maintained regardless of their eventual location choices.

Each cohort is the same size, but the cohorts' spatial distribution varies with the history of population in the economy. To make this concrete, we impose that each dying cohort of age A is replaced by the new cohort of age $a = 1$ one-for-one within their location at time of death. For example, if one quarter of the dying cohort's population lived in location j, regardless of the dying cohort's own home locations, the new cohort will have one quarter of its members with a home in location j.[33] Agents then make decisions according to the model above, with values according to (4) and choice probabilities following (6) and (8). Agents make one location decision each period.

The experiment is a regional shock that realigns the population of the economy. Specifically, the economy begins in a steady state where the locations are in two groups, with "good" or "bad" amenities. Then, at a point in time, $t = 1$, the amenities flip, good to bad and bad to good, and they remain that way indefinitely.[34] This is not the only way to envision regional disruption, but we chose it to illustrate the population transformation and migration dynamics in a way that qualitatively matches the U.S. experience in the 20th century. The newly good locations gain population share at the expense of the newly bad, and there is a transition between steady states.[35] While our implementation is artificially abrupt in order to ease the exposition, this experiment is otherwise akin to locations in the West and South "opening up" to population growth in the postwar U.S. Hence, we call the bad-to-good locations "new" and the good-to-bad locations "old."

We consider this shock in three scenarios. In the first scenario, there is no home attachment, $u_{jh} = 0$. In the second, there is a constant home attachment, $u_{jh} = \bar{u}$. In the third, the intensity of home attachment depends on the rootedness of the home location for the birth cohort, $u_{jh} = \rho R_{ja}$, where ρ is a parameter and rootedness R is the fraction of the dying cohort in location j for which j is home. In all scenarios, move costs are nonzero and constant for all cohorts and locations, $c_{oj} = \bar{c}$. We calibrate the parameters to ensure a comparable environment across scenarios; in particular, we set all to have the same steady state average migration rate.

[33] The starkness of this assumption is only for simplicity. It is not important for our broader point that an entire cohort gives birth, dies, and is replaced at their final age. It important that new cohorts' endowment of home preference be related in some fashion to their preceding cohorts' location decisions.

[34] The reason for imposing the shock in this fashion is to keep the cross-sectional variance in location quality the same before and after the shock. The cross-sectional variance in location quality affects the steady state migration rate, with more variance leading to more concentration of population and lower total migration.

[35] A period of shocks to locations followed by a period of stability does not lead to population transformation. A period of wide cross-sectional dispersion in locations followed by a period of low cross-sectional dispersion leads to an *increase* in migration because, in this model, the steady state rate of migration is *higher* when locations are more similar.

For scenario 2, we choose $\bar{c}$ and $\bar{u}$ to match in steady state the average rate of migration for the at-home and not-at-home population of 0.1 and 0.2, respectively. For scenario 3, we set $\bar{c}$ and ρ to match the same at-home/not-at-home rates as scenario 2. Finally, we set $\bar{c}$ for the first scenario (that without home attachment) to match the average rate of migration across all agents in the steady state of scenarios 2 and 3.

The plots in Figure 7 show the results of this experiment for the three scenarios. Panel A shows the economy-wide migration rate. In all scenarios, when the location quality shift occurs, migration increases from its steady state value as population reshuffles under the new spatial distribution of amenities. After that, however, the three scenarios diverge. In scenario 1 without home attachment, the migration rate falls back to its steady state value, which takes several periods because the move costs slow the adjustment process. In scenario 2 with constant home attachment, the return to steady state is more gradual because the population's relocation results in a period of transition with a more "footloose" population. This is especially stark in scenario 3 with roots-based home attachment, because the intensity of home attachment weakens during the period of transition.

Panel B shows the degree of population reallocation produced by migration, plotting the net-to-gross migration rate. Idiosyncratic shocks produce some degree of turnover among locations, with some migration flowing away from new locations and into the old locations. Moreover, scenarios 2 and 3 have horizontal location quality, in that some people prefer the old locations because of their own home attachment. This is borne out of smaller net-to-gross migration early in the transition and an elevation later on, after scenario 1 has reached steady state. In short, the population transition takes longer to obtain because agents in the model disagree about the quality of the locations.

Panels C through F break down the locations into the new and old groups; i.e., the receivers and senders of population growth in the transition period. Panel C shows the migration rate for the new locations. In our extreme experiment, migration mechanically falls as the location improves in quality and retains more of its residents. In scenario 1, this is abrupt and complete on impact. However, in scenarios 2 and 3 with home attachment, the transition period is more interesting. Though each exhibits the mechanical decrease in migration, after the initial drop, there is a phase of elevated migration as many of the newly arrived residents lack home attachment and end up moving elsewhere again, a situation reminiscent of Coen-Pirani (2010). This is especially pronounced in scenario 3 with roots-based home attachment, because the intensity of home attachment fell due to cohort parents being transplants. Panel E, below C, shows the fraction of agents in the location who are at home. In all scenarios, the at-home share falls as the population transition begins. This has no impact on migration in scenario 1, but drives the dynamics for the new locations' migration rates in panel C.

Panels D and F for the old locations are the mirror images of panels C and E, respectively.

In D, there is a mechanical increase to migration, which is immediately complete in scenario 1 but undergoes a reversal phase in the scenarios with home attachment, as shown in panel F. Those left behind in the old locations are strongly attached and less likely to migrate, a situation reminiscent of Zabek (2024).

In summary, this experiment shows the importance of home attachment for migration dynamics in an economy undergoing a regional transition. The shift towards new locations renders their populations less attached and consequently more transient, leading to a long transitional phase of elevated migration that can take many cohorts to settle. The experiment here is illustrative, however, so we proceed to estimate the impact of home attachment in the migration decision and measure its importance in the observed migration decline in the U.S.

4 Quantitative Model

In this section, we enrich the simple model in order to estimate the impact of home attachment in the individual migration decision. The empirical model extends the basic model in two ways: time and types.

4.1 Dynamic Location Choice Model

The first extension is that the decision is explicitly specified as a dynamic discrete choice. Whereas (4) has the agents making essentially myopic decisions, the agents live for multiple periods and face costs to changing locations. Hence, a recursive dynamic specification (a la Kennan and Walker (2011)) is more appropriate.[36] The choice-specific value is then

$$v_j(o) = \mu_j + u_{jh} + c_{jo} + \beta V'(j), \tag{9}$$

where $V'(j)$ is the value of beginning the next period located in j and β is a discount rate. Note that this formulation has location as a state variable (denoted in parentheses).[37] Under the two-stage nesting, the value of starting a new period in a location is the expected value of being faced with the move/stay decision:

[36]The theoretical motivation for the importance of dynamics in migration goes back to Sjaastad (1962) and Topel (1986).

[37]In general, a dynamic model would contain additional state variables such as cyclical economic conditions or personal attributes like wealth. Nothing in our framework or estimation strategy precludes states beyond location, but we do not incorporate additional state variables for several reasons. First, we are interested in long-run trends in migration and not higher frequency dynamics, so we do not include cyclical variables. Additionally, we have found that in the annual ACS data, most of the variation in migration is cross-sectional in nature; see Appendix C. Second, our data come from repeated cross sections without individual longitudinal information, and we are aggregating individuals into cells to estimate their choice probabilities. Hence, we have little ability to leverage individual state variables.

Figure 7: Migration After Regional Transformation

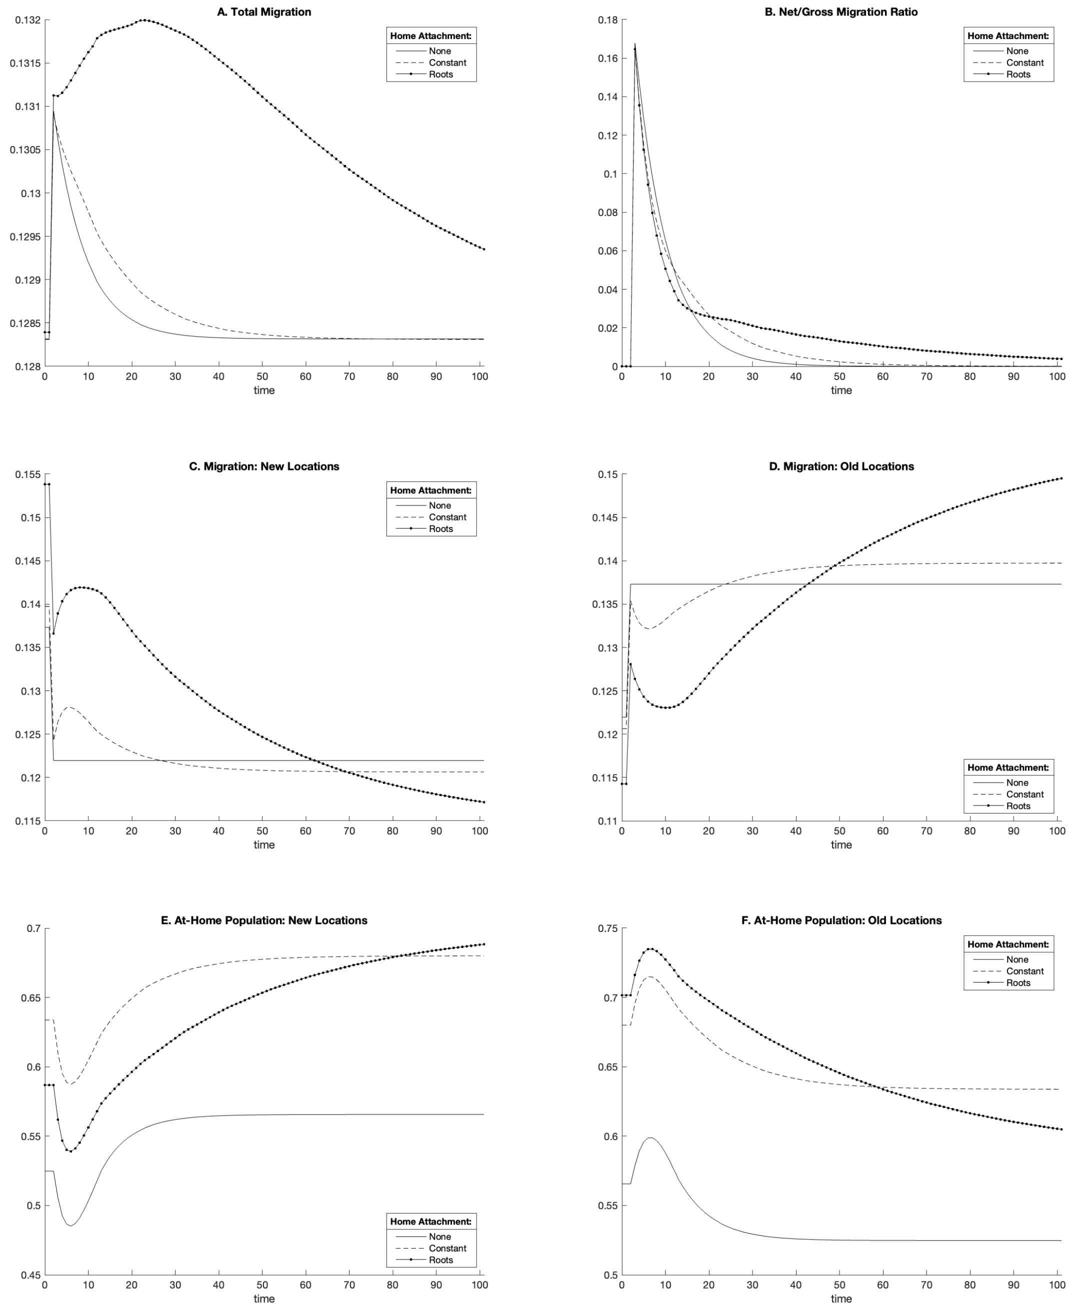

NOTES: The figures plot the timepath of variables produced by the regional shocks simulation described in Section 3.3. The shock is initiated at time $t = 1$ in the plots. Variables are denoted by their subfigure titles. (Source: Authors' calculations using simulated data.)

$$V(j) = \delta ln\big[exp[V^s(j)]^{\frac{1}{\delta}} + exp[V^m(j)]^{\frac{1}{\delta}}\big]. \tag{10}$$

Note that (9) decomposes the value of a location into flow utilities and continuation values. A key component of our estimation strategy, described below, will be to tractably account for continuation values.

4.2 Accounting for Agent Types

The second extension is to account for types. A major concern for estimation in our case is the potential difference in move propensity among types of individuals. "Types" are permanent heterogeneities and not state variables, so that the value of facing the same set of alternatives may vary by the type of individual. We alluded to this earlier in the stylized model when agents were endowed with a birthplace (but were otherwise identical).

The choice-specific values of a location are distinguished among types as

$$v_j^{\tau,H_y}(o) = u_{jh}^{\tau,H_y} + c_{jo}^{\tau} + \beta V'^{\tau,H_y}(j)\,. \tag{11}$$

In (11), τ indexes demographic groups (such as age) and H_y denotes birth type: a home location of H, born in year y. In most of the equations that follow, to ease exposition we suppress the type notation except when absolutely necessary.

This allows utility to vary demographic type in interaction with location – that is, according to whether $j = H$ and how deep the roots are of birth cohort y. Move costs vary by demographic group, but we explicitly disallow move costs to vary by home status for sake of the simulations in Section 8.

4.3 Measurement of Home Attachment

Another modeling decision is how to operationalize the concept of home attachment. We need a way to measure an abstract concept consistently. The available data allow us to use birthplace as a measure of home, as is frequently used in the literature.[38] Birthplace has the advantages of being consistently defined over time and predetermined for each individual – an endowed characteristic before they make any choices. We recognize some individuals may not sense any serious attachment to their birthplace, but its predictive power for migration shows its salience despite the possibility of measurement error.

Moreover, Section 2 also showed that the intensity of attachment among at-home residents varies between fast and slow locations, and so we introduced the concept of rootedness, the probability that one's parents were born in the same place, as a feasible proxy for the degree of

[38]See Bayer et al. (2009), Diamond (2016), Piyapromdee (2021), and Zerecero (2021).

attachment.[39] In concept, rootedness is an individual-level measure. In implementation, it is a cohort-level measure, since we cannot track the same individuals longitudinally across censuses, but must apply the average for people of the same age and birthplace. As an index, rootedness depicts the cross-geography and intergenerational differences in attachment. Rootedness has a consistent definition over time and is endowed at birth, which is useful for identifying its effect on migration. Appendix C examines the effect of our rootedness measure on migration propensity. The index is a strong predictor of out-migration among natives, but not among nonnatives. This suggests rootedness is not a place effect, but a home effect on natives. On the inflow margin, it does not appear to be a draw on return migrants.

4.4 Parameterization

We now detail the utility components from equation (11).

For home preferences, we use a linear function of roots R when the location is type H_y's home location, $u^{\tau,H_y}_{j=H} = \alpha_\tau R_{H_y} I(j = H)$, where I is the indicator function. As a specification check, we also use a simple indicator variable for home status, $\alpha_\tau I(j = H)$. In either, we allow α_τ to vary by age and education.[40]

Besides the home premium, a location affords a common flow utility available, denoted μ^τ_j, to any resident of type τ, regardless of birthplace. These control for differential attractiveness in locations in ways that can vary by group. In practice, the μ^τ_j work out to be fixed effects in estimation. Flow utility for a location is then as in (4):

$$u^{\tau,H_y}_j = \mu^\tau_j + \alpha_\tau R_{H_y} I(j = H). \tag{12}$$

Lastly, we turn to the specification of move costs. The move cost function has an intercept shifter for each type τ to account for the profile of migration over the life cycle and by worker education level. Then, to account for the spatial component of migration probability, we enter the distance in kilometers between LLM centroids. We also allow a discrete shift in distance for "neighboring" LLMs (those with counties sharing a border), for LLMs in the same state, and for LLMs in the same region. Move costs are symmetric between a pair of locations ($c^\tau_{jo} = c^\tau_{oj}$), and $c^\tau_{jj} = 0$ by construction. The move cost function is

$$c^{\tau,H_y}_{jo} = \underbrace{\sum_\tau I(\tau) c^\tau}_{types} + \underbrace{\sum_d^D c_d d_{oj}}_{distance}. \tag{13}$$

Note that individuals face moving costs whether at home or not. We thus make a distinc-

[39]The literature reviewed in Section 1.1 contains several references to theories of social network formation.

[40]α is not applicable for the foreign-born, since $I(j = H)$ is always zero for that group.

tion between moving costs and preferences for particular locations. While the home premium inclines workers to prefer their own birthplace *ceteris paribus* from any origin, moving costs introduce frictions to switching locations. These ideas are often conflated in the literature, either as shorthand or because of data limitations.[41] The distinction is especially important in counterfactual simulations, where we alter home preferences but keep moving costs fixed as a disciplining assumption.

5 Estimation Method

We next describe the estimation method. The core idea is to use the model structure to derive a set of estimating equations. We briefly describe the method here, and additional details and derivations are relegated to Appendix D.

5.1 Data Context

Before going into model details, we discuss some features of the data context to which the model will be applied. Some implementation decisions are driven by the available data and what can be identified from it.

To study a question of spatial heterogeneity in move rates, we need data as geographically rich as possible. Knowing from a long literature that demographic features such as age and education status are important for predicting migration propensity, we would like coincident demographic details as well. To study the effect of home attachment, we need information on an individual's initial location. Fortunately, census microdata fits the bill on all counts, detailing the respondent's current location, past location, birth location, and many basic individual demographic features, and covering the entire geography of the U.S. The comprehensive geography and large observation count (useful when studying a rare event such as migration) are major reasons to prefer this dataset to a small longitudinal survey.

Unfortunately, the census microdata imposes a few constraints our estimation must cope with. First, the data are cross-sectional in nature. We see each respondent once and do not observe their repeated decisions over time. Second, the only dynamic information we have is past locations (birth and last residence), and the omission of other information (such as income shocks) limits what we can do with other possible state variables. This is a tradeoff we are willing to make, because for our purposes, geography trumps individual states.

[41] For example, Moretti (2011) introduces a static model with a distribution of location-specific preferences to study their impact on labor mobility in response to local market shocks. Bayer et al. (2009), Diamond (2016), and Bryan and Morten (2019) put a measure of home in the utility function (as we do), but refer to the effects as "moving costs." Morten and Oliveira (2016) make a finer distinction, using relocation costs in their model but also checking for robustness to use of both costs and specific preferences for birthplace. Piyapromdee (2021) is similar to Diamond (2016) in execution, but terms the home effect as utility.

A final concern is that the census went from a five-year retrospective migration question in the older decennial versions of public-use microdata to a one-year retrospective in the ACS. It is not straightforward to convert between one- and five-year rates because many moves are reversed in under five years. We use one method of conversion, based on auxiliary data and some parametric assumptions, for the purpose of depicting basic empirical patterns, but we are reticent to do so in our structural model. Lacking a time series, we focus our estimation on the one-year horizon and in simulations will backcast the model onto previous periods.

5.2 Forming the Moment Conditions

Choice probabilities by type and origin are the targets of the estimating equations. With age group by education by birthplace (including foreign) types and origin states, there are $4 \times 2 \times (J+1) \times J$ cells, with J choices for each. Because spatial heterogeneity is important for our analysis, we would like to include as many locations as possible, but a larger J leads to two practical problems. First, the sample sizes for small cities become too small to reliably estimate choice probabilities. Second, because there are $J \times J$ choice probabilities and J birthplaces, the memory requirements of our stacked estimator increase cubically in J. We choose a cutoff of LLMs with at least one million residents in 2010. There are 69 named cities and a residual "outside option" location aggregating the remaining smaller places.[42] At $J = 70$, there are 39,760 types and 2,738,200 choice probabilities to estimate from the data.

Even with the large cities, however, we still encounter some small cell problems. While the move-stay decision rate is usually well-measured, some destination choices (i.e., choice probability conditional on a move) are not observed. Besides general concerns about measurement error, our estimator relies on log odds ratios, so taking the log of zero is a technical problem. We use a parametric smoothing procedure, detailed in Appendix D.[43] We first run the raw cellular data through a Pseudo Poisson Maximum Likelihood (PPML) estimator, which can handle zero observed flows, with origin-by-type and origin-destination fixed effects. The projected cell probabilities from the PPML procedure form our target moments.

5.3 The Estimation Approach

The idea behind the estimation routine is to infer model parameters from the choice probabilities (more specifically, using odds ratios). The intuition is the typical revealed-preference appeal used in the literature on the estimation of demand for heterogeneous products, which has been widely applied to location choice settings.[44] Higher choice probabilities ("market shares")

[42]The two smallest included cities are Fort Myers, Florida and Manchester, New Hampshire. The two largest excluded are Poughkeepsie, New York, and Baton Rouge, Louisiana.

[43]We are indebted to an anonymous referee for suggesting this procedure.

[44]The dynamic demand applications are the nearest analogues; see Bayer et al. (2016) or Davis et al. (2021).

imply greater demand, which implies a hierarchy of location quality. The degree to which the choice probabilities vary across types implies their relative preferences for each location, accounting for heterogeneous demand by type. Variation in types by their origins then implies the resistance between origins and destinations (i.e., distance), much like a gravity regression in the trade or transportation literatures.

Appendix D.1 contains a sketch of the revealed-preference identification argument.

5.4 Managing Dynamic Considerations

If the only objective were to predict odds ratios, the above system of equations would be sufficient. We could project odds ratios for each cell, and account for differences in the cell weights over time to simulate changes in migration due to demographics. This would be a useful exercise, but we would like to go a step further by finding the home premium utility flow implied by migration flows. Identifying the primitives, in addition to being interesting on its face, allows us more flexibility in conducting counterfactuals.

To estimate utility parameters using data on choice probabilities as in (5) and (8), one needs to account for the dynamic considerations embedded in v_j. Otherwise, estimation of utility parameters may be biased by commingling the effects of the continuation value $V(j)$ with flow utility u_j. For example, the choice of home may offer utility premia now and in the future, but originating there also changes the expected value of moving in the future relative to a non-home location. We want an estimator to account for the option value associated with this and other (possibly unobserved) features.

But observing one choice per individual, it is difficult to credibly estimate a likelihood function using a value function solution; it would involve major assumptions about how state variables will evolve for the individuals. Instead, we will account for the value of future options through the use of conditional choice probabilities (CCP), pioneered by Hotz and Miller (1993). This technique leverages the mapping between choice probabilities and values, using an estimate of choice probabilities conditional on states to stand in for the continuation value term.

CCP techniques rely on the result that the expected value of the maximum is known in closed form under T1EV preference shocks. From (3), we can substitute in for the continuation value in (4) to obtain

$$v_j(o) = u_j + c_{jo} + \beta[v_z(j) - \ln \sigma_{zj}]. \tag{14}$$

Designating some choice z as a normalization and using a pre-computed estimate of the probability of making that choice σ_{zj}, one can estimate the utility parameters in u and c without explicitly solving for $V(j)$. In our case, the choice of z and the use of $v_z(j)$ and $ln\sigma_{zj}$ is complicated by a nesting of the model, an issue which we will discuss in detail below.

5.5 Utility Parameters

For the normalization in (14), for our location choice model, we leverage the property of finite dependence (Arcidiacono and Miller (2011)) to avoid solving a dynamic programming problem in estimation. There are two parts, one for each nest of the decision model. We begin our explanation with step 2 and work backwards.

Moving: Choice of Destination

Taking from (5) the log odds ratio yields an expression for the difference of two choice probabilities as the difference in their choice values:

$$\lambda(\ln \sigma_{jo} - \ln \sigma_{ko}) = v_{jo} - v_{ko} = (u_j + c_{jo}) - (u_k + c_{ko}) + \beta(V'(j) - V'(k))\,. \tag{15}$$

The relative values of two choices contain their continuation values, which differ by the change each choice causes in the state variable, next period's location. The finite dependence property elides the computation of this object. The difference in continuation values can be substituted using (10) then (8b) to yield

$$V(j) - V(k) = (V^m(j) - \delta \ln \sigma^m(j)) - (V^m(k) - \delta \ln \sigma^m(k))\,. \tag{16}$$

With nesting, however, the value of beginning a period in a location is not defined symmetrically to arriving in the location. We need two more steps to obtain an expression with value functions differenced away. The value of moving can be substituted using (6) then (5) to obtain

$$V^m(j) = \lambda[v_{zj} - \ln \sigma_{zj}], \tag{17}$$

which expresses the value of moving from j as the choice-specific value of moving from j to a third location, z, and the probability of doing so. The values v_{zj} and v_{zk} are the same up to their moving costs, closing the loop in finite dependence:

$$v_{zj} - v_{zk} = c_{zj} - c_{zk}. \tag{18}$$

These substitutions leave an expression of (15) with only utility parameters and choice probabilities:

$$\begin{aligned}&\lambda[\ln \sigma_{jo} - \ln \sigma_{ko}] = \\ &(u_j - u_k) + (c_{jo} - c_{ko}) - \beta\delta(\ln \sigma'^m(j) - \ln \sigma'^m(k)) - \beta\lambda(\ln \sigma'_{zj} - \ln \sigma'_{zk}) + \beta\lambda(c_{zj} - c_{zk}).\end{aligned} \tag{19}$$

In comparison to a standard finite dependence, this expression contains two choice probabilities to account for continuation values – the probability of moving, and, conditional on moving, the probability of choosing a destination z – but otherwise follows the same logic of returning the problem to the same state in a future period.[45]

Moving or Staying

Another complication of the nested problem is that we need two odds ratios – one for each level of nest – so another derivation is needed for the move-or-stay level of the decision problem. From (8a) and (8b), the log odds ratio for staying versus moving is

$$\ln \sigma^s - \ln \sigma^m = \frac{1}{\delta}(V^s - V^m). \tag{20}$$

The difference in the stay/move odds ratio between two different locations is then

$$\delta[(\ln \sigma^s(j) - \ln \sigma^s(j)) - (\ln \sigma^s(z) - \ln \sigma^m(z))] = (V^s(j) - V^m(j)) - (V^s(z) - V^m(z)) \,. \tag{21}$$

This expression can be separated into moving and staying blocks. The moving block, $V^m(j) - V^m(z)$, can be derived using the procedure described above for the destination probabilities. The staying block, $V^s(j) - V^s(z)$, is dealt with by first separating it into its flow and continuation value components,

$$V_s(j) - V_s(z) = u_j - u_z + \beta(V'(j) - V'(z)) \,. \tag{22}$$

The continuation value difference returns us to (16), with z and k flipped:

$$V(j) - V(z) = \delta(\ln \sigma^m(z) - \ln \sigma^m(j)) + \lambda(\ln \sigma_{kz} - \ln \sigma_{kj}) + \lambda(c_{kj} + c_{kz}) \,. \tag{23}$$

Putting the moving and staying blocks back together, we have

$$\begin{aligned} &\delta[(\ln \sigma^s(j) - \ln \sigma^m(j)) - (\ln \sigma^s(z) - \ln \sigma^m(z))] = \\ &\qquad u_j - u_z + \beta\delta(\ln \sigma^m(z) - \ln \sigma^m(j)) + (\beta - 1)\lambda(\ln \sigma_{kz} - \ln \sigma_{kj} + (c_{kj} + c_{kz})) \,, \end{aligned} \tag{24}$$

which reflects a stability in move rates over time, conditional on type and location.

[45] Because our cross-sectional data use only location as a state variable, we assume $\sigma^m(j) = \sigma'^m(j)$ and $\sigma_{kz} = \sigma'_{kz}$. This is acceptable when the variation in move probabilities between locations and types is more important than temporal variation, which we have argued elsewhere (see Appendix C). Davis et al. (2021) take a similar tack.

Combined Estimating Equations

Combining the two nesting steps yields a set of estimating equations,

$$\lambda(\ln\sigma_{jo} - \ln\sigma_{ko}) = v_{jo} - v_{ko} \tag{25a}$$

$$\delta[(\ln\sigma^s(o) - \ln\sigma^m(o)) - (\ln\sigma^s(z) - \ln\sigma^m(z))] = (V^s(o) - V^m(o)) - (V^s(z) - V^m(z)), \tag{25b}$$

where, after applying the derivations of (19) and (24), the individual estimating equations can be stacked to form a vector equation where objects from the data can be placed on the left[46] and (functions of) parameters on the right:

$$\underbrace{\begin{bmatrix} Y_1 \\ Y_2 \end{bmatrix}}_{Y} = \underbrace{\begin{bmatrix} u(x_j) - u(x_k) & c_{jo} - c_{ko} + \beta(c_{zj} - c_{zk}) \\ u(x_j) - u(x_z) & (\beta - 1)(c_{ko} - c_{kz}) \end{bmatrix}}_{X} \underbrace{\begin{bmatrix} \theta_u \\ \theta_{mc} \end{bmatrix}}_{\theta}. \tag{26}$$

Vector Y contains combinations of choice probabilities (Y_1 from (19) and Y_2 from (24)). Matrix X is composed of functions of utilities and moving costs: for example, an indicator for whether a location is home (in the utility column) or the distance between locations (in the move costs column). Vector θ is the vector of parameters to be recovered. In short, Y is data, X is model structure, and θ are parameters. The vectors of choice probabilities can be stacked across the locations in the sample to form a set of moment conditions. Estimation proceeds much like a standard regression: The data matrix X is inverted on the choice probabilities to recover the estimand, θ.[47]

Equation (26) identifies the parameters of interest off of differences in choice probabilities. However, scale and location parameters (including β, λ, δ, and intercepts of the move costs by type) are not identified here and must be calibrated elsewhere, as we describe next.

5.6 Move Cost Intercepts

A separate estimator for the move cost intercepts derives from the absolute level of mobility within a type group, as measured by the stay/move odds ratio, controlling for the different opportunity sets offered to each origin. Move costs consist of an intercept term (for moving anywhere) and a distance term (depending on the origin-destination):

[46] The derivation of terms can be found in Appendix D.

[47] It is convenient but not strictly necessary that utility is linear in parameters. The separability of the preference shocks is required, however, to derive the mapping of value functions to CCPs. Note that an attractive feature of the regression-like specification is that utility function matrix inversion accounts for the covariances in the data. For example, if shallow-rooted places happen to be younger or less remote, the covariance between move costs and home attachment is accounted for in the X matrix.

$$c_{jo} = c_0 + c(d_{jo}).$$

Notice that because the move cost term in the value of a location j from origin o is separable, the term for the value of moving can be factored as

$$V^m(o) = \lambda \ln[\sum_{i \neq o} exp(v_{io})^{1/\lambda}] = c_0 + \lambda \ln[\sum_{i \neq o} exp(u_i + c(d_{io}) + \beta(V(i))^{1/\lambda}]. \tag{27}$$

Hence, the odds ratio of moving versus staying can be represented as

$$\delta(\ln \sigma^s - \ln \sigma^m) = V^s(o) - c_0 - \lambda \ln[\sum_{i \neq o} exp(u_i + c(d_{io}) + \beta(V(i))^{1/\lambda}]. \tag{28}$$

The moving cost intercept can be recovered from this odds ratio if one can account for the differences in the origin value $V_s, V_m - c_0$. Intuitively, this means to elicit the fundamental cost after controlling for the features that make moving more or less likely in a given origin. A derivation detailed in Appendix D.4 shows that the odds ratio can be expressed as

$$\delta(\ln \sigma^s(o) - \ln \sigma^m(o)) = c_0 + \lambda \ln(\sum_{i \neq o,z} \frac{\sigma_{iz}}{\sigma_{oz}} exp(-(c_{iz} - c_{oz}) + c(d_{io}))^{1/\lambda}). \tag{29}$$

This expression is convenient because choice probabilities are observable from the data, and the distance-dependent cost terms were recovered in the utility parameters estimation step, meaning the only remaining unknown term is c_0. Each origin location and birthplace-type cell provides an observation of the move/stay odds ratio. Constructing the term within the summation for each origin o accounts for the differential value the cell type encounters when facing a system of locations – much like a control function would account for unspecified error terms in a regression. The residual average from the comparison of move/stay odds ratios provides the estimate of the move cost intercept.[48] We recover this for each demographic group type τ.

5.7 Scale Parameters

Finally, we need to calibrate the scale parameter for each level of nest, δ for the move/stay decision and λ for the destination (conditional on moving) decision. The relationship between them is described by equation (19), comparing inflow rates to future move-out rates.

Because the utility is only identified up to scale, we cannot separately identify these parameters, but this comparison elicits their relative importance. We set $\lambda = 1$ and then seek to

[48] We also get very similar results using a more nonparametric approach, estimating (28) as a linear equation with fixed effects for origin-type groups.

estimate δ to set the scale ratio. We choose z to be a normalizing location (an outside option) so that $c_{zj} = c_{zk}$. Then, controlling non-parametrically for the two potential destinations j, k and for their pairing with origin o – i.e., a gravity-style regression with destination and origin/destination fixed effects – we can recover the δ parameter as the coefficient on the difference in move-out rates in the following regression:

$$\frac{1}{\beta}[\ln \sigma_{jo} - \ln \sigma_{ko}] + \lambda(\ln \sigma'_{zj} - \ln \sigma'_{zk}) = \frac{1}{\beta}[(u_j - u_k) + (c_{jo} - c_{ko})] + \delta(\ln \sigma'^{m}(k) - \ln \sigma'^{m}(j)). \quad (30)$$

We run this regression by demographic type τ to get scale parameters for each.

5.8 Order of Operations

In practice, the first step in estimation is the scale parameter because we need the δ parameter to set up the utility parameter estimation step. Even before that, we set β to 0.95 a priori to conform to an annual discount rate. The second step is the utility parameters estimation in section 5.5, which provides us with the main parameters of interest and the distance costs. The distance terms are then needed to recover the move cost intercepts, equation (29), in the third and final step.

6 Estimation Results

Seeking to estimate the ratio δ/λ, we fix $\lambda = 1$ arbitrarily and estimate the move/stay step scale coefficients according to (30) by demographic type. Table 3 reports the results. The parameter on scale of the upper nest is roughly 2 for all demographic groups. Because the choice probability in the upper nest depends on the value exponentiated to $1/\delta$, an estimate of $\delta > 1$ implies a lower elasticity of the migration decision to the available options outside the origin than movers have when selecting a particular destination. In particular, the move/stay decision is less affected (in percentage terms) by home status than the destination selection. The parameter estimates are slightly lower for the noncollege educated (≈ 1.7) than the college educated ($1.9 - 2.4$). The disparity implies the college educated are less sensitive (to a small degree) to external conditions when making moving decisions. In the current context, this means they are slightly less sensitive to home status and does not indicate whether they are less sensitive to intertemporal changes in regional conditions.

The scale coefficients are used in setting up the estimating equations for the utility parameter step outlined in sections 5.5 and 5.6. Table 4, presented over three parts (A-C), reports estimates of the key utility parameters and the model fits they induce. Each table part is organized into

Table 3: Parameter Estimates: Scale Parameters

Education:	Noncollege				College			
Age:	20s	30s	40s	50s	20s	30s	40s	50s
δ	1.709	1.680	1.737	1.726	2.423	1.914	2.085	1.961
	(0.002)	(0.002)	(0.002)	(0.002)	(0.004)	(0.004)	(0.003)	(0.003)

NOTES: The table reports the estimation results from the regression in equation (30). Reported standard errors treat each cell probability as an observation. (Source: Authors' calculations using ACS data.)

7 columns of specifications. The first two specifications ignore distance-based moving costs in order to set a benchmark. The next two columns (3 and 4) introduce distance-based move costs, and the following two (5 and 6) incorporate location by type fixed effects. The odd-numbered columns use a simple home indicator function to measure average home premium. The even-numbered columns use our preferred metric, roots-varying home preference. The final column (7) is a non-nested model used for comparison.

Part A reports the parameters for home utility and distance-based move cost. In column 1, with a fixed home utility and no distance-based move cost, home registers a large and significant utility bonus across all demographic types. Column 2, with a roots-based home utility, shows the same. (The coefficient on roots is slightly larger than the fixed premium because it is interacted with a variable scaled between 0 and 1, with a mean of about 0.7.) It is informative to see the ubiquity of home premia across types, but it is also difficult to interpret the magnitude of these effects given the group-specific scale parameters and lack of other coefficients in the specification. Columns 3 and 4 add distance-based move costs to the specification. Distance has a clear effect on destination selection, and there is a regional component to migration. Costs increase (i.e., become more negative) as distance between LLMs increases, although it is not completely smooth, as there are discounts apparent for LLMs in the same census region, same state, and "neighboring" LLMs (e.g., Washington, DC, and Baltimore, Maryland). Yet, notably, the regional nature of migration flows does nothing to change the home premium estimates compared with columns 1 and 2. Specifications 3 and 4 have more interpretable magnitudes, however, because home utility can be compared with distance-based move costs. Using the noncollege 20 year-olds estimate, for example, we calculate that home preference is a comparable magnitude to the preference for moving from New York to nearby Washington, DC, rather than the much-farther Dallas, Texas.

Columns 5 and 6 add location-by-type fixed effects. This flexibly controls for a host of unspecified location quality features specific to each demographic group. The home premium estimates are largely unaffected, and if anything, tend to increase slightly. Distance coefficients move slightly, which can be due to spatial correlation in location quality affecting the implied distance gradient in migration. The robustness of the home premium demonstrates that home effects (horizontal, or heterogeneous by birthplace, in nature) are distinct from location quality

Table 4: Part A. Parameter Estimates: Home Utility and Distance-Based Move Costs

		Home 1	Roots 2	Home 3	Roots 4	Home 5	Roots 6	Home 7
Home Attachment								
Noncollege	20s	0.239	0.295	0.238	0.299	0.353	0.531	1.137
		(0.003)	(0.004)	(0.002)	(0.003)	(0.002)	(0.003)	(0.013)
	30s	0.165	0.204	0.166	0.210	0.301	0.479	0.935
		(0.003)	(0.004)	(0.002)	(0.003)	(0.002)	(0.003)	(0.013)
	40s	0.113	0.132	0.119	0.144	0.273	0.414	0.924
		(0.002)	(0.003)	(0.002)	(0.003)	(0.002)	(0.003)	(0.013)
	50s	0.106	0.121	0.113	0.135	0.275	0.420	0.793
		(0.002)	(0.003)	(0.002)	(0.003)	(0.002)	(0.003)	(0.013)
College	20s	0.289	0.358	0.290	0.364	0.368	0.496	1.054
		(0.005)	(0.007)	(0.005)	(0.007)	(0.005)	(0.007)	(0.013)
	30s	0.107	0.125	0.107	0.129	0.181	0.272	0.796
		(0.004)	(0.005)	(0.003)	(0.005)	(0.003)	(0.005)	(0.013)
	40s	0.153	0.183	0.154	0.189	0.240	0.366	0.627
		(0.004)	(0.005)	(0.003)	(0.005)	(0.003)	(0.005)	(0.013)
	50s	0.158	0.188	0.160	0.195	0.275	0.404	0.612
		(0.004)	(0.005)	(0.004)	(0.005)	(0.003)	(0.005)	(0.013)
Move Cost Distance								
Noncollege	log(km)			-0.139	-0.139	-0.210	-0.224	-0.173
				(0.000)	(0.000)	(0.001)	(0.001)	(0.000)
	Same region			0.600	0.600	0.604	0.593	0.596
				(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
	Same state			0.039	0.039	0.099	0.105	0.038
				(0.002)	(0.002)	(0.003)	(0.003)	(0.003)
	Neighbor			0.731	0.731	0.702	0.675	0.750
				(0.004)	(0.004)	(0.004)	(0.004)	(0.004)
College	log(km)			-0.179	-0.179	-0.256	-0.264	-0.190
				(0.001)	(0.001)	(0.001)	(0.001)	(0.000)
	Same region			0.468	0.468	0.497	0.492	0.531
				(0.002)	(0.002)	(0.003)	(0.003)	(0.002)
	Same state			0.056	0.055	0.107	0.110	-0.091
				(0.004)	(0.004)	(0.004)	(0.004)	(0.003)
	Neighbor			0.702	0.702	0.557	0.540	0.729
				(0.005)	(0.005)	(0.006)	(0.006)	(0.004)
Specification Details								
Fixed Effects	Number	8	8	8	8	552	552	8
	Outside Option	y	y	y	y			y
	All Other J					y	y	
Move Choice:	Nested	y	y	y	y	y	y	n

NOTES: The table reports estimation results from the system of equations in (26). Each column represents a different specification. Reported standard errors treat each cell probability as an observation. (Source: Authors' calculations using ACS data.)

(vertical, or common to all, in nature), a fact elicited by the estimation identifying the home premium off of agents of the same type but who differ in their home location. The location quality terms also help to contextualize the home premium, since we can compare the home utility terms to the dispersion in location fixed effects. The ratio of home premium to the standard deviation of location quality ranges from 0.31 for college-educated people in their 30s to 0.85 for noncollege-educated people in their 20s, with an average of about 0.6. Thus, home is preferred as much as a city 3/5 of a standard deviation up the ladder in overall quality.

Column 7 runs the same regression as column 3 but without the estimated scaling coefficients from Table 3, instead imposing ($\lambda = \delta = 1$). This is the result one obtains when having organized the data the way we did, but running a standard conditional logit model. The preference parameters do not appear markedly different to the eye, owing to the strong home and distance patterns in the data. The problem with the nonnested model appears in its implied migration elasticities, an outcome that will be apparent when we discuss the baseline fit of the model specifications below.

Table 4, Part B reports the moving cost intercepts for each demographic type. These largely reflect the age profile in migration that is not otherwise picked up in the model. The estimates for the college educated appear a bit higher in magnitude, but in simulation these are scaled down by the slightly larger scale coefficients. Also, the estimates are very similar across specifications, indicating that the total amount of migration implied is not greatly affected by the specification of the opportunity set in moving value. (From the odd- to even-numbered columns with home indicators or roots preference, respectively, the estimates are the same to the decimals reported.)

Table 4, Part C shows statistics from simulations of the baseline migration rate according to the specifications in columns 1 to 7.[49] The top panel reports correlations between model and data of the LLM-level out-migration rate. All specifications are able to produce a high degree of correlation – that is, the model can replicate fast to slow locations by incorporating personal and spatial demographics. The correlation is even evident when splitting the populations by home status. The correlations are slightly better when using roots-based home attachment, predominantly among the at-home, because this better fits the pattern of lower home attachment in fast locations, as illustrated in Figure 4. The model also generates correlation, albeit to a smaller degree, in destination inflows. The correlation is naturally much better when using destination-by-type fixed effects, since these account for much more of a location's attributes (besides home status). This comes at some cost of fitting the outflow rates, however, and since outflows are our main focus, the fixed effects specifications are not our preferred specification

[49]We account for future values in the simulation by applying the choice probabilities from the mapping in (14). This means that the simulations that evaluate the estimators have the same future value component, which presses them to be more similar to each other than if they were all simulated via full value function solution. In counterfactual simulations that alter utility, described later, we update the choice probability/value function mapping because of the change to fundamentals the simulations induce.

Table 4: Part B. Parameter Estimates: Move Cost Intercepts

		Home 1	Roots 2	Home 3	Roots 4	Home 5	Roots 6	Home 7
U.S. Born								
Noncollege	20s	-7.817	-7.817	-8.021	-8.021	-8.050	-8.050	-6.636
		(0.013)	(0.013)	(0.013)	(0.013)	(0.013)	(0.013)	(0.012)
	30s	-8.840	-8.840	-9.044	-9.044	-9.073	-9.072	-7.294
		(0.013)	(0.013)	(0.013)	(0.013)	(0.013)	(0.013)	(0.012)
	40s	-9.777	-9.777	-9.975	-9.975	-10.003	-10.002	-7.752
		(0.013)	(0.013)	(0.013)	(0.013)	(0.013)	(0.013)	(0.012)
	50s	-10.174	-10.174	-10.370	-10.370	-10.397	-10.396	-8.014
		(0.013)	(0.013)	(0.013)	(0.013)	(0.013)	(0.013)	(0.012)
College	20s	-8.778	-8.778	-8.950	-8.950	-8.980	-8.980	-6.528
		(0.013)	(0.013)	(0.013)	(0.013)	(0.013)	(0.013)	(0.012)
	30s	-9.302	-9.302	-9.473	-9.473	-9.502	-9.502	-7.237
		(0.013)	(0.013)	(0.013)	(0.013)	(0.013)	(0.013)	(0.012)
	40s	-11.263	-11.263	-11.433	-11.433	-11.463	-11.463	-8.001
		(0.013)	(0.013)	(0.013)	(0.013)	(0.013)	(0.013)	(0.012)
	50s	-11.132	-11.132	-11.299	-11.299	-11.328	-11.328	-8.088
		(0.013)	(0.013)	(0.013)	(0.013)	(0.013)	(0.013)	(0.012)
Specification Details								
Fixed Effects	Number	8	8	8	8	552	552	8
	Outside Option	y	y	y	y			y
	All Other J					y	y	
Move Choice:	Nested	y	y	y	y	y	y	n

NOTES: The table reports estimation results for move cost intercepts by type from equation (29). Also estimated but not reported are coefficients for foreign-born population types. Reported standard errors treat each cell probability as an observation. (Source: Authors' calculations using ACS data.)

Table 4: Part C. Comparing Baseline Simulations

		Home 1	Roots 2	Home 3	Roots 4	Home 5	Roots 6	Home 7
LLM-level Migration Rates: Corr(model, data)								
Population:	All	0.912	0.915	0.913	0.918	0.503	0.506	0.897
	At Home	0.889	0.900	0.891	0.904	0.542	0.583	0.879
	Not at Home	0.891	0.892	0.897	0.897	0.555	0.547	0.884
	Foreign Born	0.831	0.831	0.882	0.882	0.662	0.666	0.805
	Move In	0.446	0.439	0.456	0.449	0.884	0.892	0.433
	Data							
Move-In Rates								
To Home | Moving	0.175	0.104	0.102	0.110	0.108	0.152	0.154	0.153
Move-Out Rates								
At Home	2.217	2.025	2.053	2.100	2.125	1.904	1.883	0.650
Not at Home	5.646	5.397	5.379	5.532	5.507	5.270	5.235	6.360
Foreign Born	2.800	2.446	2.447	2.475	2.476	2.237	2.230	2.057
Specification Details								
Fixed Effects	Number	8	8	8	8	552	552	8
	Outside Option	y	y	y	y			y
	All Other J					y	y	
Move Choice:	Nested	y	y	y	y	y	y	n

NOTES: The table compares simulated data from the specifications reported in Tables 4A and 4B to the actual data, with each row reporting a comparison statistic. (Source: Authors' calculations using ACS data and simulation results.)

in simulation.[50]

The remainder of the table shows the specifications' ability to match the effect of home on choice probabilities. The next panel reports conditional move-home rates; that is, the probability that a mover chooses home when it is available. Each version – including the nonnested specification – does a reasonably good job of matching the especially high propensity of migrants returning home.

The panel at the bottom reports the move-out rates for each home status. The nonnested specification misses this moment badly, implying far too low a move-out propensity for those at home, predicting a move-out rate of just 0.6 percent, when the rest of the models and the data are closer to 2 percent. Furthermore, the nonnested model over-predicts the move-out rate for those not at home. This overstated gap produces a difference between home status that is at odds with the data and undesirable for a simulation of move rates under different home attachments.

We then proceed to use our preferred specification, column 4 with roots-based attachment and distance-based move costs (but not destination fixed effects) in our baseline and counterfactual simulations. This specification best fits the cross-sectional spatial heterogeneity in move rates – the fast vs. slow locations – that are central to our study.

7 Baseline Simulation

Before attempting counterfactual simulations, we show how well the model fits the key features of the migration data as presented in Section 2. Figure 9 shows migration rates for each demographic group (four age and two education categories) by home status (foreign-born, home, and not home), beginning in the upper-left panel with an overall average of home statuses. The model is able to predict well the age profile of migration for each education group. This is largely by design, of course, since each group is given its own parameters, but not a given since the estimation procedure is not exactly identified. Importantly, the model is able to match well the difference in move rates by home status. The bars in the lower left panel, the at-home group, are well below the bars in the lower right panel of those not at home. Within each group, the age and education profile of migration is well approximated by the model.

Figure 10 is an important test of the model, its ability to replicate the fast to slow locations patterns in migration exhibited in Section 2. The figure reports the migration rates by home status (averaging across demographic groups) for fast, medium, and slow locations. Starting in the upper-left panel, we see the model is able to reproduce the rank ordering (as the correlation in Table 4C suggested), and the magnitude of the gaps between the LLM groups. Most impor-

[50]The reason is that the destination fixed effects, in matching some especially popular or unpopular destinations on the move-to equations, distort a few LLM outflow rates on the move-out side of the simulation.

Figure 8: Migration by LLM, Data and Baseline Simulation

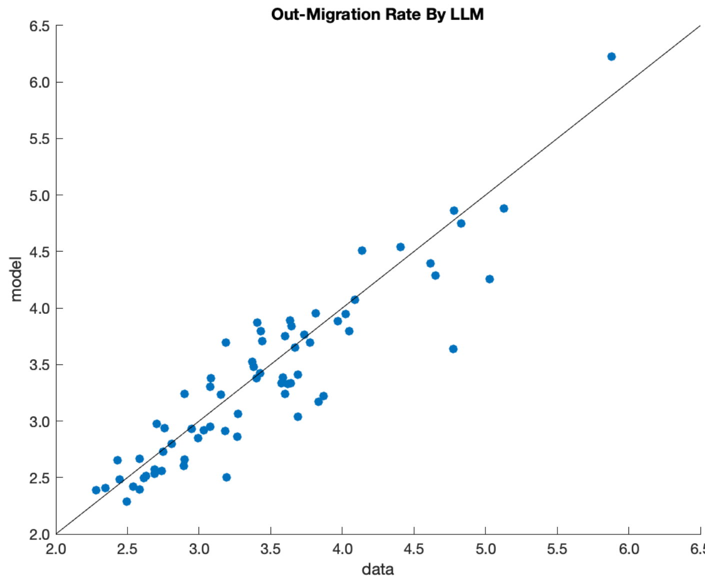

NOTES: The figure plots LLM-level migration rates from our preferred simulation (specification 4) against the actual LLM migration rates. (Source: Authors' calculations using ACS data and simulation results.)

tantly, the model exhibits the stair-step pattern from fast to slow *within* the at-home population (lower left panel), but generates no discernible pattern across LLM groups for the not-at-home, foreign- or domestic-born (righthand panels).

8 Simulation Over Time

We now turn to the major question of whether the model can produce a decline in migration over time, and critically, one that is in line with the spatial pattern in the data. The model was estimated on recent cross-sectional data, and we use the parameter estimates to simulate the migration probabilities for the population cell weighting and location attributes (chiefly, home attachment) of previous census waves. That is, we produce a backcasted time path by applying our estimated model to historical data.

8.1 Rationale of the Simulation Exercises

Our approach is a deliberate exercise to test the mechanisms in the model. The model is a structurally estimated partial equilibrium environment, which makes quantification transparent – all of the effects are direct. This is in deliberate contrast to economic geography models in which some effects are mediated through endogenous price or amenity responses.[51]

[51] See, for example, Rappaport (2004), Diamond (2016), Piyapromdee (2021), or Zabek (2024).

Figure 9: Migration by Type Group, Data and Baseline Simulation

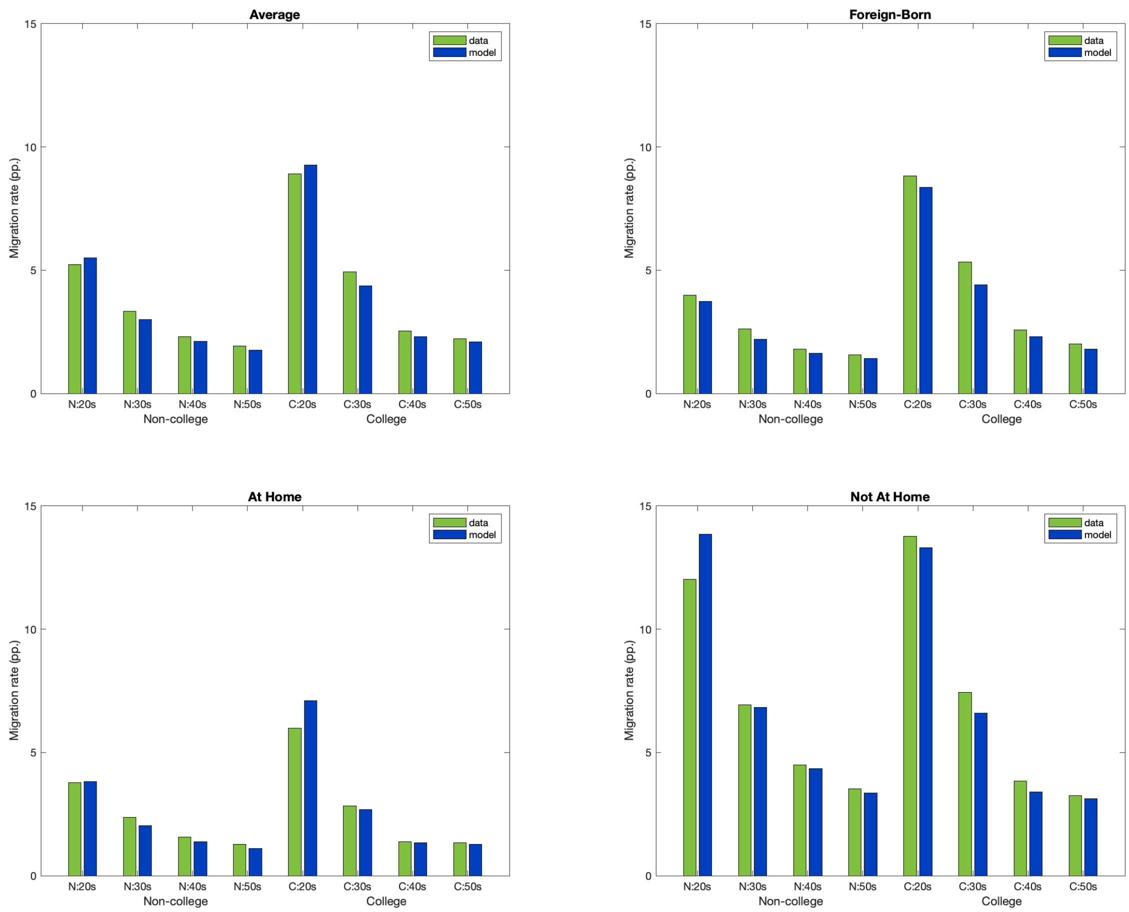

NOTES: The figure compares migration rates between our preferred simulation (specification 4) and the actual data by demographic type group and home status. (Source: Authors' calculations using ACS data and simulation results.)

Figure 10: Migration by LLM Group, Data and Baseline Simulation

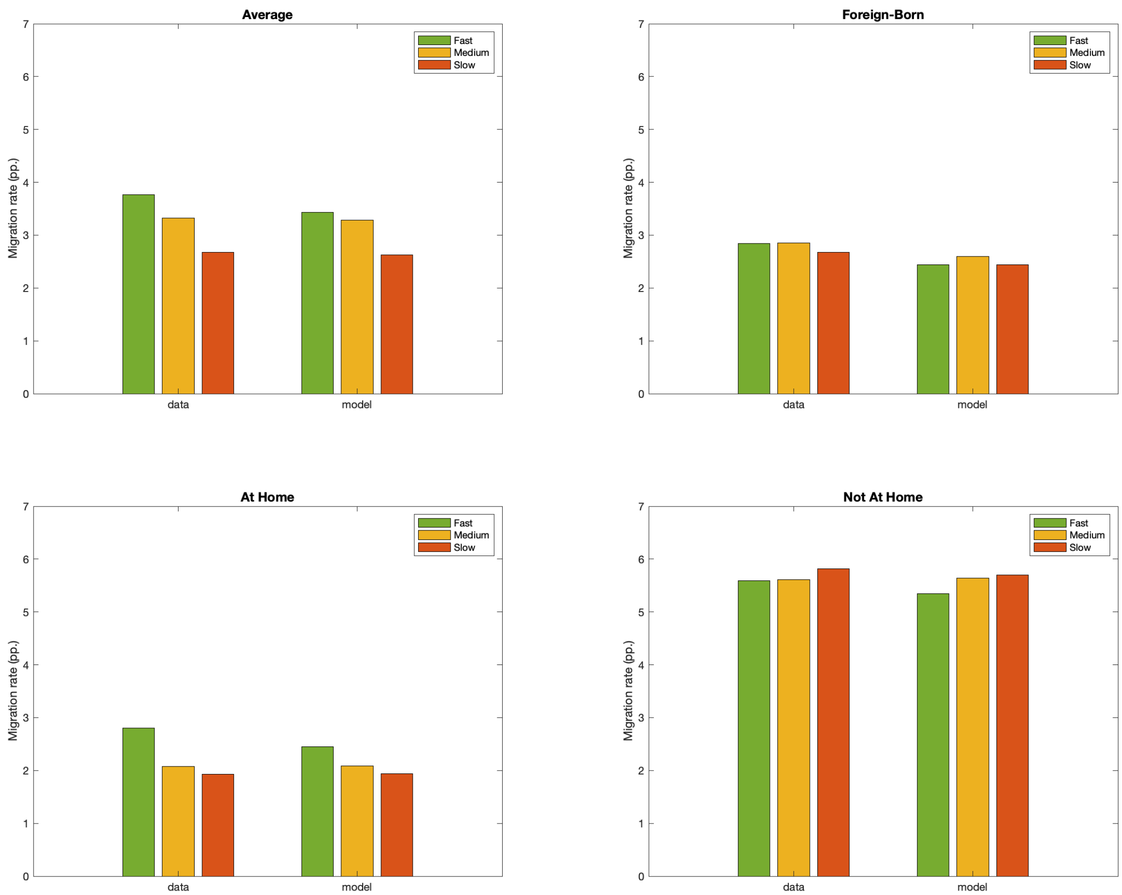

NOTES: The figure compares migration rates between our preferred simulation (specification 4) and the actual data by LLM turnover category and home status. (Source: Authors' calculations using ACS data and simulation results.)

We make some assumptions for the purpose of *ceteris paribus* hypothesis testing that we do not defend as reality. One of those assumptions is that move costs do not change. Over the timespan of data studied here, it is reasonable to believe that move costs (which represent a myriad number of personal and financial frictions) would change with technology, culture, economic conditions, and so on. But of course, a model like ours could mechanically generate a decline by assuming move costs were higher in the past or changed with distance. We would rather see if the model can replicate the decline with move costs fixed as estimated in recent data.

Another assumption is that locations did not change in overall quality. Certainly, we do not mean to defend this assumption on its face, and our own discussion in Section 2 alluded to regional transformation. But as with move costs, we want to see if the model can produce a decline without locations changing in their desirability. Moreover, our primary outcome of interest is the migration rate, which is dominated by gross (idiosyncratic) migration, while location quality has more effect on the net migration (i.e., the direction of flows that do occur). At the end, we do conduct a simulation using location quality estimates from older census waves as a check on the importance of this channel, and we find little.

With that in mind, we run our simulations as follows. We take the model's current prediction of choice probabilities. If the simulation concerns only population composition, then we take the baseline model's choice probability and simply aggregate using the cell weights for the simulated time period. For example, if we are interested in the aging of the population only, we use the baseline simulation's probabilities for each age group, and add these up using the age group weights in the respective year. When something in the environment has changed that affects utility – the main example being the rootedness of a birth cohort – we generate a new simulation of choice probabilities, solving for the continuation value by iteratively updating the conditional choice probability mapping. That is, we put the counterfactual utility values in equation (14), updating the choice probability until it reaches a steady state value.

8.2 The Combined Simulation

We begin with the model prediction from the combined simulation, with all changes to cell sizes and home utility, before decomposing the change into counterfactual simulations. Table 5 reports the migration rates for all LLMs in the estimation sample, along with their splits by fast, medium, and slow categories. The table focuses on a time period comparable to data, and then Figure 11 plots an extended time horizon all the way back to 1950 (in which, admittedly, the assumptions of fixed move costs and location quality become even more tenuous). The table shows that the model does produce a decline in mobility, with migration falling from 3.97 percent in 1980 to 3.15 in the 2010s. The decline is spatially biased as in the data, with

Table 5: Simulated Migration Rates Over Time

		1980	1990	2000	2005-2011	2012-2017
LLM Group Migration Rates						
LLMs:	All	3.972	3.751	3.360	3.173	3.154
	Slow	3.196	2.998	2.778	2.645	2.675
	Medium	4.053	3.844	3.499	3.332	3.362
	Fast	5.057	4.541	3.853	3.547	3.427
		IRS 1991-2010s	Model 1990-2005/11		Census/ACS 1990-2005/17	Model 1990-2005/17
Log Changes in Migration Rates						
LLMs:	All	0.109	0.175		0.186	0.177
	Slow	0.033	0.133		0.142	0.127
	Medium	0.093	0.152		0.197	0.146
	Fast	0.202	0.255		0.268	0.270

NOTES: The table reports the aggregate migration rates by LLM category when applying the estimates from our preferred specification (4) to current and historical census data on demographic cell type sizes and degrees of home attachment. The top panel reports the rates in percentage points, and the bottom panel compares the log changes in IRS and ACS data to the simulated data. (Source: Authors' calculations using ACS data, IRS migration data, and simulation results.)

fast locations falling by much more – 5.06 percent in 1980 to 3.43 in the 2010s. Slow and medium locations do show some decline, and the rank ordering of fast to slow locations is generally preserved across time. The lower panel of the table compares the migration change with some reliable checkpoints in the data. It reports the log change in migration from 1990 to the 2005-2011 average to compare with the IRS time series, and from 1990 to the 2005-2017 average to compare with the Census/ACS data. The model generates a drop in migration that is quantitatively comparable to the data in the aggregate and within fast, medium and slow LLM groups.

Figure 11 illustrates the timepath of the migration change. Simulated migration rates increase from 1950 to 1980 before dropping significantly into the 2000s. This wave is similar across all LLM groups and is due to changes in the age structure; especially, the Baby Boom generation moving into prime migration years and then aging out by about the year 2000. The figure also shows an especially elevated migration rate for fast locations that falls off precipitously beginning in 1980, converging towards (but not completely reaching) the national average. In all places, the trend tapers in the 2010s.

8.3 Decomposition of Channels

The model can produce a decline in migration rates even without resorting to changes in move costs or location quality. This naturally invites questions of mechanisms. In Table 6, we run counterfactual simulations to isolate the channels producing the mobility decline. By "counterfactual," we mean that we are holding fixed everything in the economy except for the

Figure 11: Simulated Migration Rate Over Time, by LLM Group

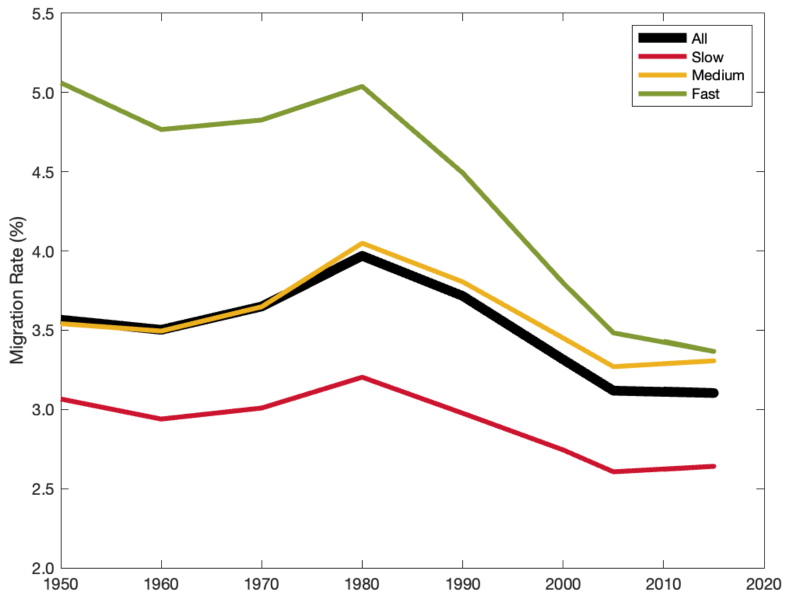

NOTES: The figure plots the timepath of migration under simulations of our preferred simulation (specification 4) by LLM category. (Source: Authors' calculations using ACS data and simulation results.)

channel(s) in question, and the economy is moving as if only the dimension of focus has changed over time. For example, to study the effect of rootedness alone, we keep all cell sizes and location features the same as in the baseline simulation, but walk back the degree of rootedness for each birth cohort as in the data. To study a channel that affects population cell size, we project a counterfactual cell size using the baseline and the ratio of the dimension in focus. That is, if the cell size for location j, age a, education e and birthplace h is

$$n_{jaeh} = n_j \omega_{aeh},$$

i.e., the product of total LLM size n_j and fraction of LLM population in the cell ω_{aeh}, we project a change in the age structure between year t and year 0 (the baseline) as

$$n^f_{jaeh} = \frac{\omega^t_a}{\omega^0_a} n_j \omega_{aeh}.$$

A change in LLM population (without a change in the composition of that population) is

$$n^f_{jaeh} = \frac{n^t_j}{n^0_j} \omega_{aeh},$$

and so on for other simulations. In some cases, we combine two or more changes. Note that nonlinearities mean two single-variable decompositions need not sum to their joint effect. Table

6 presents the decompositions as the changes between 1980 and the 2012-2017 average. The total difference in migration rate (from Table 5) is reported in the top row.

Personal Demographics. The counterfactual simulations begin in row 2 with the aging of the population referenced in the discussion of Figure 11. As suggested by the previous figure, aging does factor in significantly to the total decline in mobility. Aging alone would have produced a 0.65 percentage point decline in mobility, a sizable share of the total projected decline. However, aging was offset by another trend, the rising educational attainment of the population. Row 3 shows education upheld the migration rate by 0.15 percentage points in our model's estimate. Together, the cross-trends of aging and rising educational attainment produced a decline of 0.40 percentage point (row 4). Notably, the demographic effects are fairly similar across LLM groups. The effects are slightly larger in percentage points in fast locations, but this is simply applying a similar elasticity to a higher base rate.

Hence, demographics are important, but not the only story. We do not mean to diminish the relevance of aging or other demographic changes in the migration decline. Rather, our focus followed the literature, which right from the outset (Molloy et al. (2011)) was more interested in economic mechanisms that made the decline puzzlingly pervasive across so many demographic groups.

Row 5 projects mobility as if the population sizes across locations evolved as in the data, but within LLMs, population composition remained unchanged. Fast locations are still growing more quickly than slow locations, which has partially offset the migration decline by about 0.10 percentage point.

Row 6, "Combined Demographics," projects mobility if the population had evolved as in the data, but there were no changes within locations to the birthplace composition or the rootedness of the birth cohorts. The other factors would have resulted in a 0.28 percentage point decline in migration, with some tilt towards fast locations due to their higher base rate of migration.

Spatial Demographics. The next panel focuses on home attachment, first altering the birthplace population composition of the population (holding fixed other demographics), then the rootedness of the birth cohorts, and finally combing the two. Birthplace composition alone (row 7) would produce a 0.20 percentage point decline in mobility, predominantly centered in fast locations. This is the convergence in the at-home share exhibited in Figure 6. Rising rootedness of the cohorts (row 8) augments this effect, adding a 0.13 percentage point to the decline in fast locations, but approximately zero in others, for a total contribution of an additional 0.06 percentage point. Together, home attachment trends make for a 0.22 percentage point decline in migration in the aggregate, coming almost entirely from fast locations (row 9).

Thus, home attachment rivals personal demographic changes (specifically, aging) in its effect on the migration decline. Importantly, home attachment drives much of the spatial pattern

Table 6: Decomposition of Migration Change, 1980 to 2005/17

(in percentage points)

		LLM Group		
Simulation	All	Slow	Medium	Fast
1. Combined	0.866	0.562	0.743	1.673
Personal Demographics				
2. Age	0.653	0.530	0.689	0.738
3. Education	-0.155	-0.181	-0.196	-0.103
4. Age + Education	0.408	0.258	0.405	0.543
5. Location Size	-0.104	-0.014	-0.010	-0.036
6. Combined Demographics	0.280	0.240	0.407	0.508
Spatial Demographics				
7. Birthplace	0.190	0.052	0.055	0.403
8. Roots	0.056	0.034	-0.036	0.138
9. Birthplace + Roots	0.219	0.080	-0.011	0.498
10. Demographics + Home Attachment	1.024	0.540	0.731	1.651
Location Quality Dispersion				
11. Location Quality (1990)	0.047	0.135	0.028	-0.023
12. Location Quality (2000)	0.035	0.107	-0.016	0.004

NOTES: The table reports the change in migration under various counterfactual simulations as denoted by row. The measured outcome is decline, so a positive number indicates a decrease in the rate of migration. (Source: Authors' calculations using ACS data and simulation results.)

of the decline in fast locations. Moreover, these effects are not actually "rivals," but potentially complements. Rising home attachment can interact with and amplify other demographic changes' effects on migration. This can be seen by considering that the two subtotal rows, 6 and 9, do not sum to the total effect in row 1. To emphasize this point, row 10 simulates the changes of personal demographics and home attachments jointly to see their combined effect. In this simulation, we hold fixed population sizes across LLMs but change the population cell composition and rootedness of the cohorts. This combination would have produced an even larger migration decline, a full percentage point in the combined total, with a strong spatial bias tilting towards fast locations.

Figure E3 displays plots of the subtotal counterfactual simulations by LLM category. Each category shows similar bumps in migration due to the Baby Boom generation passing through their mobile years in the decades 1970-1990. But fast LLMs also show a downward trend from home attachment that is not present in slow or medium LLMs. The age and home attachment effects then interact to produce the larger decline evident in fast locations.

In summary, all of the considered changes in the population have some effect on migration, but without incorporating home attachment, we cannot account for the size or spatial pattern of the migration decline.

Location Quality. As a final exercise, we check whether changes in location quality have had an effect on migration rates. To do so, we estimate the model using census data from 1990 and 2000 (which contain 5-year migration rates) and draw out the location fixed effects implied by those rates. The location-by-type fixed effects capture the net location attractiveness as implied by location choices. We conduct simulation using the older years' location-by-type fixed effects with recent years' population cell sizes and home attachments.

We find the fixed effects are highly correlated across estimation samples, with a correlation coefficient of 0.7 to 0.85 depending on the demographic group, and there is not an obvious trend in the dispersion of location quality. These estimates suggest limited scope for location quality convergence itself in affecting migration rates in the last three decades. The final two rows of Table 6 report the change in migration from simulations fixing personal and spatial demographics but altering location quality according to our estimates. Directionally, these simulations do produce a decline, but it is quantitatively small, and it is centered among slow locations, not fast. This result indicates that locational heterogeneity may be converging and contributing somewhat to a migration decline, but its contribution is minor compared with demographics and home attachments.

9 Conclusion: Has America Lost Its Mojo?

This paper has presented a new explanation for the decline in geographic mobility in the U.S., starting with the novel fact that fast locations, the destinations of population growth for much of the 20th century, drove the decline. We show the decline is explainable by increased rates of home attachments in these places after population growth rates converged – which happened before migration declined.

We introduce a new angle for the study of labor and population dynamics. Our findings demonstrate the importance of spatial demographics for aggregate population processes in addition to individual propensities. When people have attachment to past locations, perhaps even across generations, population history affects population dynamics in the future.

Our result has mixed implications for policy and research. The first is simply to clarify what is and is not indicated by the mobility decline. Contrary to popular belief, mobility is declining in places that have been and are still growing relative to others, not in distressed, "left behind" areas. Moreover, the direct welfare implications of our results are actually quite sanguine. The mobility "friction" we find results from people making advantageous choices to stay in place.

Yet, we also show that concern is not unfounded. Horizontal preferences for particular locations (though not technically "move costs") will reduce responsiveness to shocks, both at an individual and population level. Rising home attachment may be a headwind to the labor market in that it reduces elasticities to local economic shocks. Further research is warranted to understand labor dynamics – gross flows and the net changes they produce – in a more spatially tied economy, including the potential role for place-based policy. Zabek (2024) is a good step in this direction.

Another implication is that the population convergence itself must now be evaluated. Our findings indicate the migration decline is a consequence of the spatial population transition of the 20th century, a period of American history with new technologies meeting sparsely populated lands. In one sense, the "phase" interpretation is reassuring – there is nothing particularly wrong with people today. But if desirable population changes are not occurring because access to good locations is cut off, whether by policy or market power, the original concerns about misallocation are still relevant, if not for the reasons first expected.

References

Alesina, A., Y. Algan, P. Cahuc, and P. Giuliano (2015): "Family Values and the Regulation of Labor," *Journal of the European Economic Association*, 13, 599–630.

Arcidiacono, P. and R. A. Miller (2011): "Conditional Choice Probability Estimation of Dynamic Discrete Choice Models with Unobserved Heterogeneity," *Econometrica*, 79, 1823–1867.

Artuc, E., S. Chaudhuri, and J. McLaren (2010): "Trade Shocks and Labor Adjustment: A Structural Empirical Approach," *American Economic Review*, 100, 1008–45.

Autor, D. H., D. Dorn, and G. H. Hanson (2013): "The China Syndrome: Local Labor Market Effects of Import Competition in the United States," *American Economic Review*, 103, 2121–68.

Barnett, C. (2008): *Mirage: Florida and the Vanishing Water of the Eastern US*, University of Michigan Press.

Baum-Snow, N. and L. Han (2024): "The Microgeography of Housing Supply," *Journal of Political Economy*, 132, 1897–1946.

Bayer, P., N. Keohane, and C. Timmins (2009): "Migration and Hedonic Valuation: The Case of Air Quality," *Journal of Environmental Economics and Management*, 58, 1–14.

Bayer, P., R. McMillan, A. Murphy, and C. Timmins (2016): "A Dynamic Model of Demand for Houses and Neighborhoods," *Econometrica*, 84, 893–942.

Bayoumi, T. and J. Barkema (2019): "Stranded! How Rising Inequality Suppressed U.S. Migration and Hurt Those Left Behind," Tech. rep., International Monetary Fund.

Belot, M. and J. Ermisch (2009): "Friendship Ties and Geographical Mobility: Evidence from Great Britain," *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 172, 427–442.

Bishop, K. C. (2008): "A Dynamic Model of Location Choice and Hedonic Valuation," Olin School of Business, Washington University in Saint Louis.

Blanchard, O. J. and L. F. Katz (1992): "Regional Evolutions," *Brookings Papers on Economic Activity*, 1, 1–37.

Bound, J. and H. J. Holzer (2000): "Demand Shifts, Population Adjustments, and Labor Market Outcomes During the 1980s," *Journal of Labor Economics*, 18, 20–54.

Bryan, G. and M. Morten (2019): "The Aggregate Productivity Effects of Internal Migration: Evidence from Indonesia," *Journal of Political Economy*, 127, 000–000.

Büchel, K., M. V. Ehrlich, D. Puga, and E. Viladecans-Marsal (2020): "Calling from the Outside: The Role of Networks in Residential Mobility," *Journal of Urban Economics*, 119, 103277.

Cadena, B. C. and B. K. Kovak (2016): "Immigrants Equilibrate Local Labor Markets: Evidence from the Great Recession," *American Economic Journal: Applied Economics*, 8, 257–290.

Caliendo, L., M. Dvorkin, and F. Parro (2019): "Trade and Labor Market Dynamics: General Equilibrium Analysis of the China Trade Shock," *Econometrica*, 87, 741–835.

Carrington, W. J. (1996): "The Alaskan Labor Market During the Pipeline Era," *Journal of Political Economy*, 104, 186–218.

Chinitz, B. (1986): "The Regional Transformation of the American Economy," *American Economic Review*, 76, 300–304.

Clark, W. A. V. and W. Lisowski (2019): "Extending the Human Capital Model of Migration: The Role of Risk, Place, and Social Capital in the Migration Decision," *Population, Space and Place*, 25, e2225.

Coen-Pirani, D. (2010): "Understanding Gross Workers Flows Across U.S. States," *Journal of Monetary Economics*, 57, 769–784.

Cooke, T. J. (2011): "It Is Not Just the Economy: Declining Migration and the Rise of Secular Rootedness," *Population, Space and Place*, 17, 193–203.

——— (2013): "Internal Migration in Decline," *The Professional Geographer*, 65, 664–675.

Correia, S., P. Guimarães, and T. Zylkin (2020): "Fast Poisson Estimation with High-Dimensional Fixed Effects," *The Stata Journal: Promoting Communications on Statistics and Stata*, 20, 95–115.

Couture, V., C. Gaubert, J. Handbury, and E. Hurst (2024): "Income Growth and the Distributional Effects of Urban Spatial Sorting," *Review of Economic Studies*, 91, 858–898.

Dahl, M. S. and O. Sorenson (2010): "The Migration of Technical Workers," *Journal of Urban Economics*, 67, 33–45.

Dao, M., D. Furceri, and P. Loungani (2017): "Regional Labor Market Adjustment in the United States: Trend and Cycle," *Review of Economics and Statistics*, 99, 243–257.

DaVanzo, J. (1983): "Repeat Migration in the United States: Who Moves Back and Who Moves On?" *Review of Economics and Statistics*, 65, 552–559.

David, Q., A. Janiak, and E. Wasmer (2010): "Local Social Capital and Geographical Mobility," *Journal of Urban Economics*, 68, 191–204.

Davis, M. A., J. Gregory, D. A. Hartley, and K. T. Tan (2021): "Neighborhood Effects and Housing Vouchers," *Quantitative Economics*, 12, 1307–1346.

Davis, S. J., R. J. Faberman, and J. Haltiwanger (2012): "Labor Market Flows in the Cross Section and Over Time," *Journal of Monetary Economics*, 59, 1–18.

Dawkins, C. J. (2006): "Are Social Networks the Ties That Bind Families to Neighborhoods?" *Housing Studies*, 21, 867–881.

Decker, R. A., J. Haltiwanger, R. S. Jarmin, and J. Miranda (2016): "Declining Business Dynamism: What We Know and the Way Forward," *American Economic Review*, 106, 203–07.

DeWaard, J., M. Hauer, E. Fussell, K. J. Curtis, S. D. Whitaker, K. McConnell, K. Price, D. Egan-Robertson, M. Soto, and C. A. Castro (2022): "User beware: Concerning Findings from the Post 2011–2012 US Internal Revenue Service Migration Data," *Population Research and Policy Review*, 41, 437–448.

Diamond, R. (2016): "The Determinants and Welfare Implications of US Workers' Diverging Location Choices by Skill: 1980-2000," *American Economic Review*, 106, 479–524.

Falck, O., S. Heblich, A. Lameli, and J. Südekum (2012): "Dialects, Cultural Identity, and Economic Exchange," *Journal of Urban Economics*, 72, 225–239.

Fischer, C. S. (2002): "Ever-More Rooted Americans," *City & Community*, 1, 177–198.

Frey, W. (2009): "The Great American Migration Slowdown," Tech. rep., Brookings Institution, Washington, DC.

Ganong, P. and D. Shoag (2017): "Why Has Regional Income Convergence in the US Declined?" *Journal of Urban Economics*, 102, 76–90.

Glaeser, E. L., J. Gyourko, and R. E. Saks (2006): "Urban Growth and Housing Supply," *Journal of Economic Geography*, 6, 71–89.

Glaeser, E. L., D. Laibson, and B. Sacerdote (2002): "An Economic Approach to Social Capital," *Economic Journal*, 112, F437–F458.

Gyourko, J., C. Mayer, and T. Sinai (2013): "Superstar Cities," *American Economic Journal: Applied Economics*, 5, 167–199.

Heise, S. and T. Porzio (2019): "Spatial Wage Gaps and Frictional Labor Markets," Tech. Rep. 898.

Herkenhoff, K. F., L. E. Ohanian, and E. C. Prescott (2018): "Tarnishing the Golden and Empire States: Land-Use Restrictions and the US Economic Slowdown," *Journal of Monetary Economics*, 93, 89–109.

Hotz, V. J. and R. A. Miller (1993): "Conditional Choice Probabilities and the Estimation of Dynamic Models," *Review of Economic Studies*, 60, 497–529.

Hsieh, C.-T. and E. Moretti (2019): "Housing Constraints and Spatial Misallocation," *American Economic Journal: Macroeconomics*, 11, 1–39.

Hyatt, H., E. McEntarfer, K. Ueda, and A. Zhang (2018): "Interstate Migration and Employer-to-Employer Transitions in the United States: New Evidence from Administrative Records Data," *Demography*, 55, 2161–2180.

Hyatt, H. R. and J. R. Spletzer (2013): "The Recent Decline in Employment Dynamics," *IZA Journal of Labor Economics*, 2, 5.

Institute for Social Research (2021): "Panel Study of Income Dynamics, Public Use Dataset," Tech. rep., University of Michigan.

Jovanovic, B. (1979): "Job Matching and the Theory of Turnover," *Journal of Political Economy*, 87, 972–990.

Kan, K. (2007): "Residential Mobility and Social Capital," *Journal of Urban Economics*, 61, 436–457.

Kaplan, G. and S. Schulhofer-Wohl (2017): "Understanding the Long-Run Decline in Interstate Migration," *International Economic Review*, 58, 57–94.

Karahan, F. and S. Rhee (2014): "Population Aging, Migration Spillovers, and the Decline in Interstate Migration," FRB of New York Staff Report, no. 699.

Kennan, J. and J. R. Walker (2011): "The Effect of Expected Income on Individual Migration Decisions," *Econometrica*, 79, 211–251.

Koşar, G., T. Ransom, and W. Van der Klaauw (2021): "Understanding Migration Aversion Using Elicited Counterfactual Choice Probabilities," *Journal of Econometrics*, 231, 123–47.

Lkhagvasuren, D. (2012): "Big Locational Unemployment Differences Despite High Labor Mobility," *Journal of Monetary Economics*, 59, 798–814.

Luckingham, B. (1984): "The American Southwest: An Urban View," *Western Historical Quarterly*, 15, 261–280.

Ma, L. (2019): "Learning in a Hedonic Framework: Valuing Brownfield Remediation," *International Economic Review*, 60, 1355–1387.

Manson, S., J. Schroeder, D. V. Riper, and S. Ruggles. (2018): "IPUMS National Historical Geographic Information System: Version 13.0 [Database]," Tech. rep., Minneapolis: University of Minnesota.

Michielin, F., C. H. Mulder, and A. Zorlu (2008): "Distance to Parents and Geographical Mobility," *Population, Space, and Place*, 14, 327–345.

Molloy, R., C. L. Smith, and A. Wozniak (2011): "Internal Migration in the United States," *Journal of Economic Perspectives*, 25, 173–196.

——— (2017): "Job Changing and the Decline in Long-Distance Migration in the United States," *Demography*, 54, 631–653.

Molloy, R., R. Trezzi, C. L. Smith, and A. Wozniak (2016): "Understanding Declining Fluidity in the US Labor Market," *Brookings Papers on Economic Activity*, 2016, 183–259.

Monras, J. (2018): "Economic Shocks and Internal Migration," Tech. rep., CEPR Discussion Paper 12977.

Moretti, E. (2011): "Local Labor Markets," *Handbook of Labor Economics*, 4, 1237–1313.

——— (2012): "What Workers Lose by Staying Put," *Wall Street Journal*, May 26.

Morten, M. and J. Oliveira (2016): "Paving the Way to Development: Costly Migration and Labor Market Integration," Working Paper 22158, National Bureau of Economic Research.

Mulder, C. H. and G. Malmberg (2014): "Local Ties and Family Migration," *Environment and Planning A: Economy and Place*, 46, 2195–2211.

Munshi, K. (2003): "Networks in the Modern Economy: Mexican Migrants in the U.S. Labor Market," *The Quarterly Journal of Economics*, 118, 549–599.

Notowidigdo, M. J. (2011): "The Incidence of Local Labor Demand Shocks," Tech. rep., National Bureau of Economic Research.

Partridge, M. D., D. S. Rickman, M. R. Olfert, and K. Ali (2012): "Dwindling U.S. Internal Migration: Evidence of Spatial Equilibrium or Structural Shifts in Local Labor Markets?" *Regional Science and Urban Economics*, 42, 375–388.

Piyapromdee, S. (2021): "The Impact of Immigration on Wages, Internal Migration, and Welfare," *Review of Economic Studies*, 88, 406–453.

Rappaport, J. (2004): "Why Are Population Flows So Persistent?" *Journal of Urban Economics*, 56, 554–580.

Ravenstein, E. G. (1885): "The Laws of Migration," *Journal of the Statistical Society of London*, 48, 167–235.

Reisner, M. (1993): *Cadillac Desert: The American West and Its Disappearing Water*, Penguin.

Ruggles, S., S. Flood, R. Goeken, J. Grover, E. Meyer, J. Pacas, and M. Sobek (2019): "Integrated Public Use Microdata Series: Version 9.0 [Machine-Readable Database]," Tech. rep., Minneapolis: University of Minnesota.

Saks, R. E. (2008): "Job Creation and Housing construction: Constraints on Metropolitan Area Employment Growth," *Journal of Urban Economics*, 64, 178–195.

Sjaastad, L. A. (1962): "The Costs and Returns of Human Migration," *Journal of Political Economy*, 70, 80–93.

Thompson, D. (2016): "How America Lost Its Mojo," *The Atlantic*, May 27.

Topel, R. (1986): "Local Labor Markets," *Journal of Political Economy*, 94, 111–143.

Trippett, F. (1979): "The Great American Cooling Machine," *Time Magazine*, 13.

Tsivanidis, N. (2023): "Evaluating the Impact of Urban Transit Infrastructure: Evidence from Bogota's Transmilenio," Working paper.

Wiggins, L. (1995): "The Birth of the City of Miami," *Tequesta-The Journal of the Historical Association of Southern Florida*, 5–38.

Yagan, D. (2019): "Employment Hysteresis from the Great Recession," *Journal of Political Economy*, 127, 2505–2558.

Zabek, M. (2024): "Local Ties in Spatial Equilibrium," *American Economic Journal: Macroeconomics*, 16, 287–317.

Zabel, J. E. (2012): "Migration, Housing Market, and Labor Market Responses to Employment Shocks," *Journal of Urban Economics*, 72, 267–284.

Zerecero, M. (2021): "The Birthplace Premium," Working paper.

Appendices

A Details on Data Construction

The following section contains additional details on the processing procedures for datasets used in this paper.

A.1 Sources

Our migration data come from two sources, the American Community Survey (ACS) and the migration flows tables from the U.S. Treasury's Internal Revenue Service. The ACS, from Integrated Public Use Microdata Sample (IPUMS), reports the respondent's current and one-year-ago Public Use Microdata Area (PUMA) of residence, from which we can elicit migration probability (move or not) and direction (origin-destination pairs). We use the ACS from 2005 to 2019. Migration is elicited using the *puma* and *migpuma* variables. For the destination choice probability used in the imputations for moment conditions (described in Sections 5.2 and D.6), we use aggregated flow tables of ACS flows. We process the 1990 and 2000 census data in the same way. These datasets are similar to the ACS, but the retrospective window for the migration inference question ("Where did you reside t years ago?") is five years instead of one, as described above.

The IRS infers migration events from changes in the address on individual tax returns in two successive years, publishing the total county-to-county flows in each year, as well as the total stayers in, inflows to, and outflows from individual counties. One limitation is that the data are censored at flows less than 10 households, meaning many origin-destination pairs are unobserved. On average, about 70 percent of flows are on observed origin-destination routes, and the rest are censored. We measure migration using the internal subtotals of total domestic inflows and outflows, subtracting flows between counties within the same local labor market (which are rarely censored).

The IRS data underwent a change in method in the 2011-2012 tax year that resulted in noticeable differences in the sample represented. DeWaard et al. (2022) caution "user beware" in analyzing IRS migration data after 2011. Our understanding is that the data were computed and published by the Census Bureau from 1990 to 2011, and the IRS took charge in 2012 and following. The IRS has different methods for tracking addresses across multiple returns (in cases of, for instance, household formation and dissolution), and late filers, which tend to be households with complicated returns. Thus, the set of individuals represented changed, and because of the recursive nature of the data, this introduced year-over-year fluctuations that may take several more years before they can be safely compared across time. We present the data for the period 2012-2019, but rely only on the consistent sample of 1990-2011.

We also leverage aggregated population data at the county level, which we will use to show population growth trends. We obtained the county population estimates from National Historical Geographic Information System and relied heavily on that project's harmonization of geographies across census years. Census microdata samples from 1880-2000, used in the calculations of roots, home status, and income distributions, were also obtained from IPUMS. Intercensal year population estimates were obtained from the U.S. Census Bureau.

Information on location and income dynamics comes from the Panel Study of Income Dynamics (PSID). Estimates of relative income position by state of residence were obtained by merging reported PSID labor income with state-level income distributions by year, education, and foreign-born status obtained from the Annual Social and Economic Supplement of the Current Population Survey. All incomes are deflated by the consumer price index from the Bureau of Labor Statistics.

A.2 Geography: The Local Labor Market

In this paper, we will work with a unit of analysis we term a local labor market (LLM), which fully partitions the geography of the continental U.S. The LLM is derived from a commuting zone (CZ) but modified to meet some specific objectives. One objective is geographic consistency over time and across datasets. We were able to define constant boundary LLMs for both counties and PUMAs for use in, respectively, Census aggregate population and microdata, dating from 1880 to current releases. A second objective is to fit more intuitive notions of an integrated labor market area, more like a core-based statistical area (CBSA) or metropolitan statistical area (MSA). In many cases, these line up well with the commuting zone, but in some, the CZ covers a large and heterogeneous area. For example, most of southern California is in one CZ, despite substantial heterogeneity in populations and labor market opportunities between the inland counties, which we split into a Riverside/San Bernardino LLM, and the coastal counties, which we further split into Los Angeles and Ventura LLMs, making three local labor market units instead of one.

In our descriptive analyses in section 2, we report measures of migration flow and population growth for the 183 LLMs that are characterized as urban areas.[52] We aggregate the remainder of the continental 48 states into an omitted category, comprising rural areas and some unusual LLMs – smaller cities dominated by universities ("college towns" such as Athens, Georgia, or Bloomington, Indiana) or military bases (such as Jacksonville, North Carolina), which have non-standard migration behavior. Our empirical model focused on the 69 largest LLMs, aggregating the remainder LLMs into the residual location.

[52]The working paper version contains a full list of LLMs, and a dataset of the mapping of counties and PUMAs over time is available on our webpage.

Throughout, we define a migration event as an exit from an LLM for a different LLM, so that a move within a county or PUMA, or across counties or PUMAs within the same LLM, is considered staying in place. Distances between LLMs were calculated using Great Circle distance from central county population centroids for the U.S. Census.

A.3 Discussion of the Calculation of Roots

The decennial census data contain geographic information on current residence and birth state for all individuals in a household. We use household structure variables and cohort matching to estimate rootedness of a particular cohort for each home LLM. We identify a birth cohort by looking at all individuals who are less than 10 years old in a particular census wave. For example, a twenty-something in 2010 was aged less than 10 in 1990. For the cohort living in each LLM, we calculate the percentage of their parents who were born in a state in which the LLM has a county. For example, children in Dallas are rooted if their parents were born in Texas, and children in Kansas City are rooted if their parents were born in Kansas or Missouri. The child must be living in his/her home LLM to be counted in this sample. We ascribe the cohort-LLM combination to have the rootedness measured by this fraction.

There are a few possible concerns given that we use LLM for location definition but state for place of birth. We do not actually know in which *city* a child's parents were born. It is possible that a child was born in Dallas but the parents were born in Houston or Austin, though we assign them all as at home in Texas. When comparing across cohorts, if the measurement error is similar, the *change* in rootedness is still accurate. But if we are comparing an LLM in a state with several large cities, such as Texas, to an LLM in a state with only one major city, such as Minnesota, we will likely measure the rootedness of Dallas as too high compared to Minneapolis. This may be grounds for some within-state migration adjustment. In practice we found that adjustments made trivial impact on our rootedness measures, because out-of-state birthplaces drove the first-order differences between cities. Unless within-state migration is strongly negatively correlated with between-state (which other datasets indicate is not the case), our measure of rootedness will if anything shrink the dispersion of rootedness across LLMs.

We use the location of residence for children under 10 as "home" for the purpose of cohort matching. There is some mobility of young children that introduces uncertainty into our estimates. One possible adjustment is to probabilistically assign children to potential birth cities, but in practice the change to measurements is small. (As described separately, we do assign a probabilistic weight for birth LLM within a state of birth when designating types, but not in the calculation of roots.) Note that by the time we have matched cohorts, we have divided our adult sample into college graduates and non-college graduates. Since college graduates are more mobile on average, and there is positive intergenerational transmission of education, we expect that the college graduate subset of any cohort will be less rooted than the cohort as a whole. Our methodology implicitly assigns the relationship between rootedness and education to be the same in every city.

A final note is that the rootedness proportions only include the U.S. born. The foreign born are treated as a separate birthplace group – not at home, and not from other U.S. – since, among other concerns, they have no domestic "home" location to prefer above others, and they cannot be rooted. However, another concern is first-generation immigrants. In such cases, the child was born in the United States but one or both parents were born outside the U.S., leaving ambiguity in defining the child's roots. By our strict definition, of course we can say with certainty this child was not born in the same commuting zone as his parents. However, many immigrants move to cities that have an established population of immigrants from their native country already.[53] For our purposes, it was simplest to calculate rootedness only for children of native-born parents, since only they can be categorized as "at home," and since the correlation of immigrant share to total LLM turnover is weak.

A.4 Conversion of One-Year to Five-Year Migration Rate

The census and ACS microdata have the great advantage of containing demographic information including home status. Unfortunately, the census changed from a five-year retrospective question (asking where the household lived five years ago) in the decadal form to a one-year question in the ACS, making it difficult to draw comparisons of migration rates over time, because there is good reason to believe the N-year rate is not an N-multiple of the one-year rate. Fairly often, moves are reversed (return to origin) or repeated (moved again, to a new location); see DaVanzo (1983) and Kennan and Walker (2011). Hence, some one-year moves will not be observed at a five-year horizon. Moreover, because of the state dependence exhibited in return and onward decisions – recent movers move again – the five-year transition matrix is not equivalent to the one-year transition to the fifth order.

We use auxiliary longitudinal survey data with information on move frequency by demographic type, home status, and move history to convert one-year rates to implied five-year rates. The conversion uses the following formula:

$$\sigma_{5yr} = 5 \cdot \sigma_{1yr}[1 - p_{reverse} - 0.5 \cdot p_{onward})].$$

The one-year arrival rate of a move is reduced by (i) the probability of a reversal ($p_{reverse}$) occurring within five years, meaning the move would not be observed at all in a survey, and (ii) the probability of another move event to a new location (p_{onward}) occurring within five years,

[53] Munshi (2003) and Piyapromdee (2021) discuss migration networks of immigrants.

Figure B1: Changes in LLM Migration Rates

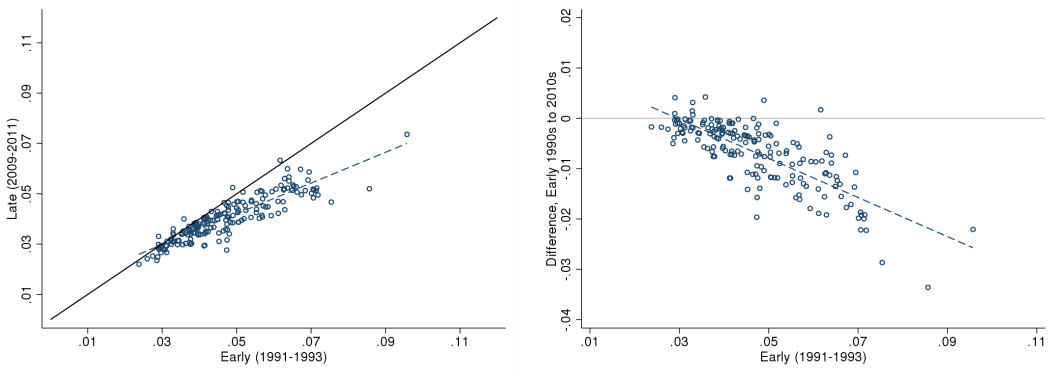

NOTES: The left figure plots migration rate of LLMs in the early 1990s to their rates in the early 2010s. The right figure plots the migration rate in the early 1990s to the difference between the rates in early 2010s and the early 1990s. The figures use 3-year averages. (Source: IRS data.)

meaning only one of the two moves would be observed. These adjustments deflate the projection to something that would match a five-year retrospective question.

What remains is to find estimates of the reversal and onward probabilities. For longitudinal data detailing the location history of individuals, we turn to the PSID. This provides a long record of individuals' location histories in addition to demographic information. For each move in the data, we can observe whether the individual moved back to the original or onward to a third location, and if either occurred, the time elapsed since the first move. This allows us to estimate a probability of reversal and onward migration within a t-year window, including estimates by age, education, and home status at the time of move. Our preferred specification uses a survival time model measuring the probability that the initial move "survives" to the fifth year as a function of age, education, and home status at time of the initial move, although findings are robust to other specifications.

B Additional Results on LLM Mobility Rates

In this appendix, we provide some additional results describing the migration patterns across U.S. LLMs. The figures presented here are scatterplots of the data that is summarized into fast, medium, and slow categories in Section 2.

Figure B1 shows the relationship between initial migration (circa 1990) and the decline over the next two decades; this figure can be compared with Figure 2. The lefthand plot displays early migration rates against late. It shows a clear off-diagonal tilt, indicating the faster the LLM, the more it slowed. The righthand plot uses the same data and takes the difference, emphasizing that while there are few markets increasing in migration rates, the declines are much larger in faster LLMs.

Figure B2 uses ACS data to show the relationship between LLM migration rates and the

Figure B2: Migration Rate and the At-Home Share

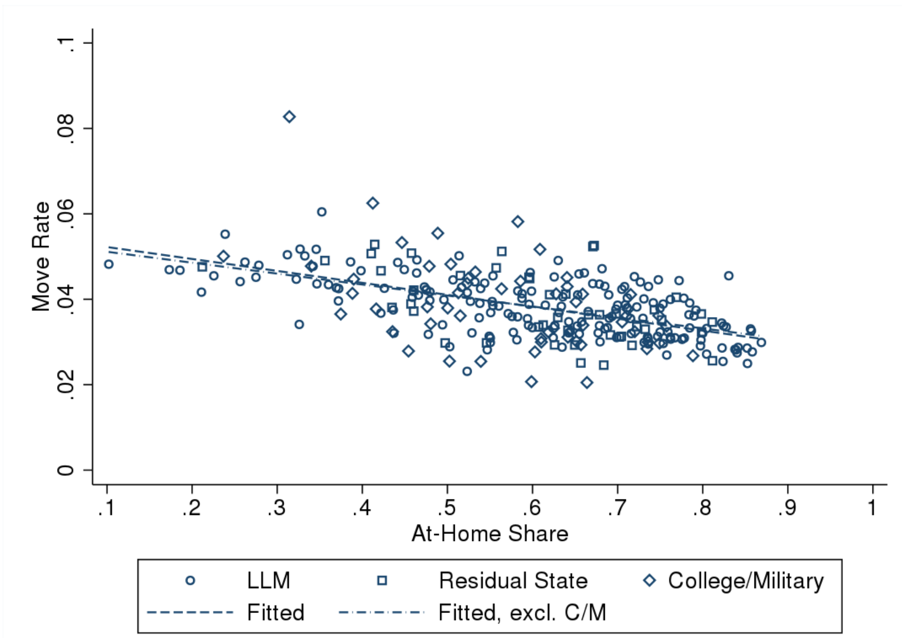

NOTES: The figure plots migration rates of LLMs against the share of their population residing at home. Plot marks distinguish between normal metro areas (LLMs), rural areas grouped into residual U.S. states, and small metro areas with a large share of university or military employment ("College/Military" towns). (Source: ACS data.)

share residing at home; this figure can be compared with Figure 3. There is some variance among the LLMs, but there is a clear downward trend throughout the support of the data. All the LLMs are included in this figures, with plot mark symbols denoting some places excluded from our main analysis. These are rural counties not in metro areas, collected into residual states, and small cities with a high share of college or military occupational employment. These do not upset the main pattern.

Figure B3 plots migration propensity against the degree of rootedness of the LLM population. Since rootedness varies by cohort, each dot is an age/education/LLM cell, and then we need to normalize migration propensity to compare across types. The data on the vertical axis is the log difference between the migration rate of the at-home to those not at home in the same age/education type by LLM cell. The plot shows that greater rootedness implies lower migration of the at-home cohort relative to the same type, not at home.

Figure B4 plots the decline in migration rates of LLMs against the share of their populations at home. This relationship has some dispersion, but on average, a lower at-home share indicates greater declines. The LLMs with much of their population at home show little to no decline in migration, while the large decliners are LLMs with low at-home share.

Figure B3: Migration Rate of the At-Home and Rootedness

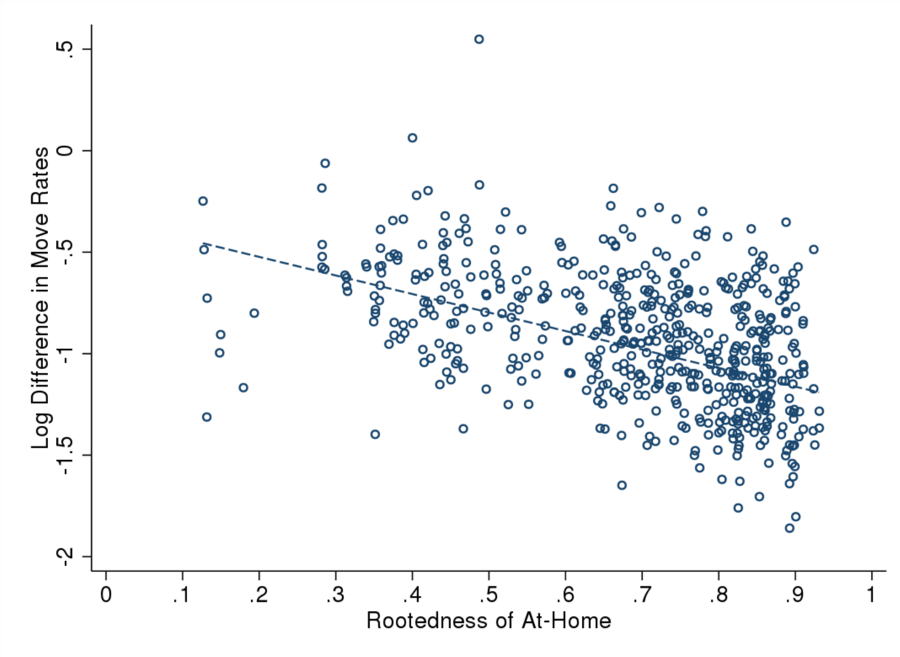

NOTES: The figure plots the migration rate log difference between an at-home cohort type and the same type, not at home, in the same LLM. A dot corresponds to an age/education/LLM cell. (Source: ACS data.)

Figure B4: Changes in LLM Migration Rates and the At-Home Share

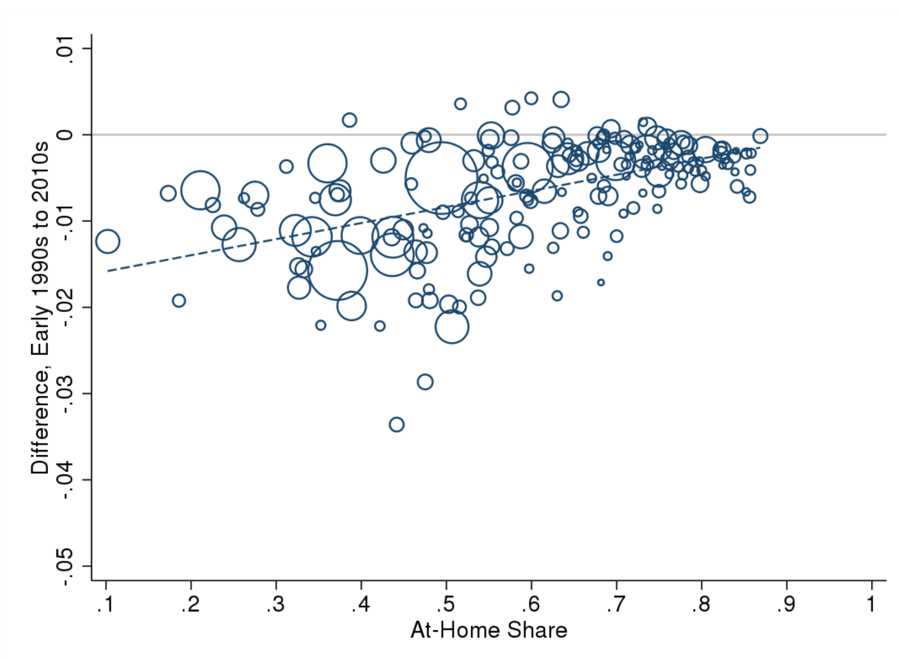

NOTES: The figure plots the change in LLM migration rate (as in Figure B1) against the at-home share of the population (as in Figure B2). (Source: ACS and IRS data.)

Table C1: R^2 From Regressions of Migration Probability on Cell Attributes

	1	2	3	4	5	6	7	8
R2	0.181	0.086	0.039	0.040	0.045	0.277	0.289	0.294
Fixed Effects								
Age X Education	8							
Birthplace Status		3						
LLM			70	70			70	
Year				13			13	
Age X Educ. X Birthplace Status						24	24	24
LLM X Year					907			907

NOTES: (Source: ACS microdata.)

C Additional Results on Migration Propensity

In this appendix, we provide some background results on the migration behavior across LLMs. Throughout, we focus on ACS microdata.

C.1 Type, Space, and Cyclical Factors in Migration

In Table C1, we conduct an exercise to quantify the amount of variation in migration that can be attributed to several categorical factors. The outcome variable is the migration probability in a cell group, where cells are stratified by age, education, home status, LLM, and year of observation. We then run a linear regression of the move probability by cell on the categorical variables that interact to comprise the cells. The table reports the R^2 from these regressions to quantify the amount of variation due to the dimensions included in the regression.

Column 1 includes the eight type categories by age and education (decade bin by college status). These alone describe 18 percent of the migration probability. Column 2 uses only home status (home, not home, and foreign born). These three categories alone explain almost 9 percent. By comparison, columns 3 and 4 show LLM and year of observation categories (70 + 13), which account for just 4 percent of variation. Interacting LLM by year in column 5 only marginally increases the explained variance. When we interact age, education and home status in column 6, the explained variance rises to 28 percent from 18 in column 1. Layering LLM, year, or LLM by year on top only slightly increases the explained variance.

Our takeaway is that type categories are important for describing the frequency of migration, more so than the location and time factors. Certainly it is important to understand the cyclical factors and trends that underly regional evolutions, but in terms of measuring the likelihood of migration, types are paramount. This is why we focus on type categories – especially home status – when quantifying the long-run trends in gross migration rates, despite this being an unusual focus relative to much of the economics literature on migration.

C.2 Odds Ratio Analysis: Reconciling Inflow and Outflow Elasticities

Table 2 showed (i) that persons at home were less likely to move away from their current location than those not at home, and (ii) that persons away from home were more likely to choose home when migrating. Qualitatively, the home biases point in the same direction, but quantitatively, there is a large asymmetry between inflow and outflow – differences simply too large to be reconciled by the conditional logit model.

To see the issue, consider a simple odds ratio rendering of the logit model. A person with home h resides in origin o. The odds of choosing a location j relative to another location k is $\hat{\sigma}^h_{jk} = \frac{\sigma^h_{j,o}}{\sigma^h_{k,o}}$. Using (3) and taking logs, the relative odds ratio is

$$\ln \hat{\sigma}_{jk} = v_{j,o} - v_{k,o}$$

Assume for the sake of illustration that locations are identical except for the home match premia, denoted ϕ, offered to their natives. We will compare outflow to inflow odds ratios.

First compare a move-out from home with a move-out from another origin not home; $k = o = h$, $j \neq h$. Then

$$\ln \hat{\sigma}_{jk} = v_{j,o} - v_{k,o} = -\phi.$$

Now consider for an agent not at home, $o \neq h$, a move-in to home to a move-in to another non-home alternative; $k \neq o$ and $j = h$. Then

$$\ln \hat{\sigma}_{jk} = v_{j,o} - v_{k,o} = \phi.$$

Hence, the two directions are symmetric, with the outflow odds ratio being the inverse of the inflow ratio.

But such symmetry is at odds with the data. Table C2 reports odds ratio regressions with separate specifications for each age by education type group. While there are small but noticeable differences in the magnitudes across groups, all share a marked asymmetry between in- and out-migration on the effect of home status, with the move-in odds ratio far exceeding the inverse of the move-out ratio. Thus, the need for a flexible elasticity between in- and out-margins is apparent for all groups.

Thus, the standard conditional logit model cannot account for the asymmetry in sensitivity to home location when deciding whether to move out and deciding where to move in. This is not entirely surprising, as other research (see Monras (2018)) has shown the asymmetric elasticities of in- and outflow migration to local labor market attributes. Following Monras (2018), we employ a nested version of the logit choice model, which segments the decision into a

Table C2: Odds Ratio Regressions of Moves Out and Moves In On Birthplace Measures

	Noncollege 20s	Noncollege 30s	Noncollege 40s	Noncollege 50s	College 20s	College 30s	College 40s	College 50s
Move Out	0.339	0.365	0.371	0.368	0.448	0.390	0.394	0.439
	(0.006)	(0.007)	(0.008)	(0.008)	(0.012)	(0.012)	(0.013)	(0.015)
Year FE	x	x	x	x	x	x	x	x
N Individuals	2,688,426	3,203,806	4,136,521	4,844,458	641,743	1,632,035	1,763,353	1,832,659
Choice	0/1	0/1	0/1	0/1	0/1	0/1	0/1	0/1
Move In	21.588	17.487	16.359	14.462	17.328	13.748	10.312	9.813
	(0.793)	(0.757)	(0.751)	(0.671)	(0.864)	(0.71)	(0.616)	(0.61)
Distance	x	x	x	x	x	x	x	x
N Individuals	41,118	30,465	27,470	27,562	23,930	31,612	18,110	17,086
Choice	1 of 69	1 of 69	1 of 69	1 of 69	1 of 69	1 of 69	1 of 69	1 of 69

NOTES: Each coefficient reports the change in the relative odds of a migration event (out of or in to a location) by the characteristic indicated in the row. A coefficient of one indicates no effect; less than one a reduction in the odds, and greater than one an increase in odds. All explanatory variables are defined as indicators for the characteristic defined in the row. (Source: ACS microdata.)

move/stay stage, and conditional on moving, a "where-to" stage of selecting a new destination. This distinction is a modeling device to permit the elasticity of choice with respect to location attributes to vary between outflow and inflow stages.

C.3 Roots: A Person or Place Attribute?

In concept, rootedness is a personal attribute – whether one's own family is born in the same place. In practice, we applied a cohort-level matched average. In the recent migration data, we do not actually know where an individual's parents were born, but we can associate the average for same-age, same-birthplace individuals from prior census waves. Using a proxy variable naturally raises questions of measurement error. Additionally, rootedness may be correlated with other place features that affect migration behavior (among other things) for all people residing in that place.

Tables C3 and C4 examine the effect of home status – and rootedness in particular – on migration propensity when controlling for other LLM-level attributes. Here we switch to linear probability models to interpret marginal effects and to accommodate a large number of fixed effects.

Table C3 reports on the effect of home status in the move-out decision probability. Being at home (column 1) and higher rootedness when at home (column 2) reduce this likelihood. Within at-home status, moreover, greater rootedness leads to even lower migration probability (column 3). Notably, while this effect is negative for the at-home, it is slightly positive for those not at home, who migrate away from more rooted places at higher rates (column 4). This indicates rootedness is an at-home effect, not a general place effect.

Column 5 shows these results are robust to inclusion of local labor market features. Moves out are less likely in places with higher mean incomes and more likely in places with more dispersion, but the results on home status are unaffected. Columns 6 and 7 add fixed effects, LLM and year (in 6) and LLM by year (in 7). These control non-parametrically for attributes of the LLMs that lead to higher or lower migration rates on average. The effects of home status are similar, with the coefficient on rootedness for the at-home increasing in magnitude. In specifications with LLM fixed effects, more identifying variation comes from between-cohort differences, indicating that within a home origin, more rooted generations are even more reluctant to migrate away. The coefficient on roots for those not at home flips sign, indicating that rootedness among the at-home is not necessarily a disamenity for nonnatives, but the attributes of the LLMs that correlate with roots lead to more out-migration.

Table C4 then studies the marginal effects of roots on destination selection (moves in). First, it expands the sample to all movers, not only those away from home. The effect of rootedness on move-in propensity can then be decomposed into all possible home statuses – a move to

Table C3: Linear Probability Regressions of Moves Out on Birthplace Measures and LLM Attributes

	1	2	3	4	5	6	7
Home	-0.037		-0.027	-0.022	-0.020	-0.021	-0.021
	(0.001)		(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Roots X Home		-0.046	-0.013	-0.013	-0.018	-0.040	-0.040
		(0.001)	(0.003)	(0.003)	(0.003)	(0.004)	(0.004)
Roots X Not Home				0.007	0.004	-0.018	-0.019
				(0.002)	(0.002)	(0.004)	(0.004)
Mean(Income)					-0.032		
					(0.002)		
Var(Income)					0.024		
					(0.006)		
Age X Education X Foreign-born Categories	x	x	x	x	x	x	x
Year FE	x	x	x	x	x	x	
LLM FE						x	
LLM x Year FE							x

NOTES: Each coefficient reports the change in the relative odds of a migration event (out of or in to a location) by the characteristic indicated in the row. A coefficient of one indicates no effect, less than one a reduction in the odds, and greater than one an increase in odds. All explanatory variables are defined as indicators for the characteristic defined in the row. (Source: ACS microdata.)

home, a move elsewhere when home was an available destination, or a move elsewhere when home was not available (either because the person was at home in her origin or because the person is foreign born). Columns 1 through 3 run through these possible specifications. The results show that home is the destination with outsized importance in the choice set, but greater rootedness does not increase the likelihood of coming home. However, rootedness decreases the likelihood of in-migration from all groups, not just potential returnees to home, suggesting greater rootedness is present in less desirable LLMs on average. When we control for local incomes (columns 4 and 5), location pair distance nonparametrically (column 5), the negative effect of rootedness declines in magnitude. It declines further for the not-at-home samples as we layer in LLM, year, and LLM by year fixed effects in columns 6 through 9. All of these results, with the reminder that more rooted places tend to be growing more slowly in population, are consistent with rootedness being higher in less desirable locations, though rootedness itself can be an amenity for natives of the location.

In summary, we find that rootedness is a good predictor of migration behavior despite its being a proxy variable for a more nebulous concept. The fact that it predicts migration for the at-home but not others suggests it picks up the social and/or psychological attachments we have in mind but cannot measure directly.

Table C4: Linear Probability Regressions of Moves In on Birthplace Measures and LLM Attributes

	1	2	3	4	5	6	7	8	9
Home	0.1179	0.1184	0.1183	0.1176	0.1151	0.1132	0.1121	0.1132	0.1121
	(0.0008)	(0.0008)	(0.0008)	(0.0008)	(0.0008)	(0.0008)	(0.0008)	(0.0008)	(0.0008)
Roots X Home		-0.0483	-0.0483	-0.0479	-0.0418	-0.0340	-0.0374	-0.0340	-0.0373
		(0.0054)	(0.0054)	(0.0054)	(0.0054)	(0.0054)	(0.0054)	(0.0054)	(0.0054)
Roots X Not Home, Not Home Sample			-0.0055	-0.0049	0.0060	0.0103	0.0134	0.0104	0.0135
			(0.0002)	(0.0002)	(0.0003)	(0.0004)	(0.0004)	(0.0004)	(0.0004)
Roots X Not Home, At Home Sample			-0.0380	-0.0370	-0.0309	-0.0223	-0.0233	-0.0222	-0.0232
			(0.0003)	(0.0003)	(0.0003)	(0.0005)	(0.0005)	(0.0005)	(0.0005)
Roots X Not Home, Foreign Born Sample			-0.0092	-0.0081	0.0029	0.0065	0.0101	0.0065	0.0101
			(0.0002)	(0.0002)	(0.0003)	(0.0004)	(0.0004)	(0.0004)	(0.0004)
Mean(Income)				0.0064	0.0043				
				(0.0001)	(0.0001)				
Var(Income)				0.0136	0.0005				
				(0.0007)	(0.0007)				
Constant	0.0160	0.0160	0.0160	-0.0620	-0.0351	0.0198	0.0140	0.0198	0.0140
	(0)	(0)	(0)	(0.0013)	(0.0014)	(0)	(0)	(0)	(0)
Distance	x	x	x	x		x		x	
Orig-Dest Pair FE					x		x		x
LLM FE						x	x		
LLM x Year FE								x	x

NOTES: Each coefficient reports the change in the relative odds of a migration event (out of or into a location) by the characteristic indicated in the row. A coefficient of one indicates no effect, less than one a reduction in the odds, and greater than one an increase in odds. All explanatory variables are defined as indicators for the characteristic defined in the row. (Source: ACS microdata.)

D Model Estimation Details

This appendix provides more detailed information on the estimation procedure. It begins with an explanation of the intuition behind identification. Then it provides the derivations of the model's estimation equations.

D.1 The Estimation Approach

The estimation approach is to use type by origin by destination choice probabilities in a system of equations relating the model parameters to choice outcomes. This approach is motivated in large part by the data available: cross-sectional data with a large number of locations and types, but no longitudinal information on individuals.

D.1.1 Using Flows to Infer Model Parameters

To show how migration flow probabilities identify the preference parameters, we start with an example. Consider a simple two-location economy with a symmetric moving cost and one type of agent. The complete list of origin-destination value functions is as

$$\begin{aligned} v_{11} &= \mu_1 & v_{21} &= \mu_2 + c \\ v_{12} &= \mu_1 + c & v_{22} &= \mu_2, \end{aligned}$$

where c is the move cost and μ is a (generic) location quality. With information on choice probabilities, we can identify the location quality terms and move costs via the difference in stay/move odds ratios between the two locations:

$$\begin{aligned} \ln\sigma_{11} - \ln\sigma_{21} &= \mu_1 - \mu_2 - c \\ \ln\sigma_{21} - \ln\sigma_{22} &= \mu_1 - \mu_2 + c. \end{aligned}$$

This has two equations and two unknowns (with the model identified only up to scale). Intuitively, the difference in location quality is identified by the net flows, and the move cost is identified by the size of gross flows relative to net.

This is the essential idea behind our identification approach. We set up a series of estimating equations, with (functions of) choice probabilities on the lefthand side and (functions of) parameters on the right. The parameters are then identified via a large regression model. To see this, consider the values of alternative locations j, k, relative to a normalizing choice z, out of origins n, o. This yields us four equations,

$$\begin{aligned} v_{jn} - v_{zn} =& (\xi_j + c_{jn}) - (\xi_z + c_{zn}) \\ v_{kn} - v_{zn} =& (\xi_k + c_{kn}) - (\xi_z + c_{zn}) \\ v_{jo} - v_{zo} =& (\xi_j + c_{jo}) - (\xi_z + c_{zo}) \\ v_{ko} - v_{zo} =& (\xi_k + c_{ko}) - (\xi_z + c_{zo}), \end{aligned}$$

where $\xi_j = u(j) + \beta V'(j)$ is shorthand for the net present value of the destination, and we are (for now) suppressing agent-type subscripts.

Carrying forward this algebra yields a $J \times (J-1)$ system of equations with the location values as destination fixed effects and the move costs as pairwise distance functions, much like a gravity equation. Because the log odds ratio of choosing a destination j is $\ln \sigma^m_{jo} - \ln \sigma^m_{zo} = v_{jo} - v_{zo}$, we can use relative choice probabilities to identify the average location quality and the resistance imposed by distance between locations.

The system illustrated above is for the choice probability of a destination, conditional on moving. Our nested model introduces an additional complication in that move/stay probabilities are distinct from (but not independent of) destination choices. However, a similar odds ratio system equations applies:

$$\begin{aligned} \ln \sigma^s_j - \ln \sigma^m_j = V^s_j - V^m_j =& \xi_j - \lambda ln\Big(\sum_{i \neq o} exp(v_i(j))^{\frac{1}{\lambda}}\Big) \\ \ln \sigma^s_k - \ln \sigma^m_k = V^s_k - V^m_k =& \xi_k - \lambda ln\Big(\sum_{i \neq o} exp(v_i(k))^{\frac{1}{\lambda}}\Big). \end{aligned}$$

The intuition is the same as before: relatively higher staying/moving odds ratios imply relatively higher location qualities. Connecting with the equations above, this set of equations merely substitutes in for the diagonal entries in the matrix of migration flows. It is somewhat more involved (algebraically) to account for the expected value of moving from one origin compared with another, but we will describe this below in the discussion about accounting for continuation values.

D.1.2 Accounting for Types

Before moving on, we show how the system of equations by type helps identify the preferences by type – including, notably, preference for home. While the geography of the system provided variation that allowed us to identify the relative values of ξ_j vis-a-vis ξ_k, it is the variation between types that identifies the home premium and type-specific quality valuations.

For example, consider a type τ with a home of location $H = j$ compared with another type τ with a home location of $H = k$:

$$(v_{jn}^{\tau,H=j} - v_{zn}^{\tau,H=j}) - (v_{jn}^{\tau,H=k} - v_{zn}^{\tau,H=k}) = \xi_j^{\tau,H=j} - \xi_k^{\tau,H=k} = \alpha_\tau R_{H_y} + \beta V_{t+1}^{\tau,H=j}(j) - \beta V_{t+1}^{\tau,H\neq j}(j). \tag{31}$$

The fact that these agents of the same type τ have the same valuations – but for the home premium of j – identifies how large the additional premium is via any difference in choice odds ratios.

A similar reasoning applies for identifying the valuation differences between different types. For example, consider a type $\tau = 1$ compared with a type $\tau = 2$, each with the same home of location h.

$$(v_{jn}^{\tau=1,h} - v_{zn}^{\tau=1,h}) - (v_{jn}^{\tau=2,h} - v_{zn}^{\tau=2,h}) = \xi_j^{\tau=1,h} - \xi_j^{\tau=2,h} = \mu_j^{\tau=1} - \mu_j^{\tau=2} + \beta V^{\tau=1,h}(j) - V^{\tau=2,h}(j). \tag{32}$$

A difference in odds ratio therefore implies one type values the location more than the other.

D.2 Model Equations

To begin the derivations, we collect the model equations together. Throughout, we suppress type subscripts for legibility sake.

The value of moving into a destination j from an origin o is

$$v_{jo} = u_j + c_{jo} + \beta V'(j). \tag{33}$$

For movers out of location o, the probability of choosing destination j is

$$\sigma_{jo} = \frac{exp(v_{jo})^{1/\lambda}}{\sum_{i\neq o} exp(v_{io})^{1/\lambda}}. \tag{34}$$

The probability of staying in the origin is

$$\sigma_s = \frac{exp(V^s)^{1/\delta}}{exp(V^s)^{1/\delta} + exp(V^m)^{1/\delta}}. \tag{35}$$

The probability of moving to a new destination is

$$\sigma_m = \frac{exp(V^m)^{1/\delta}}{exp(V^s)^{1/\delta} + exp(V^m)^{1/\delta}}. \tag{36}$$

The ex ante value of a location is

$$V(o) = \delta \ln[exp(V^s)^{1/\delta} + exp(V^m)^{1/\delta}]. \tag{37}$$

The value of remaining in a location (i.e., the value of staying or not moving) is

$$V^s = u_o + \beta V'(o). \tag{38}$$

The value of moving out of a location (i.e., the expected value of the maximum destination value) is

$$V^m = \lambda \ln[\sum_{i \neq o} exp(v_{io})^{1/\lambda}]. \tag{39}$$

D.3 Utility Parameters

There are two forms of estimating equations in the nested model: moving/staying choice and the destination conditional on moving choice. We begin with the latter.

Destination choices

From (34), the difference in log moving probabilities between two destination options is

$$\ln \sigma_{jo} - \ln \sigma_{ko} = \frac{1}{\lambda}(v_{jo} - v_{ko}) = \frac{1}{\lambda}[(u_j - u_k) + (c_{jo} - c_{ko}) + \beta(V'(j) - V'(k))]. \tag{40}$$

The difference in continuation values can be substituted using (37) then (36) to get

$$V(j) - V(k) = (V^m(j) - \delta \ln \sigma^m(j)) - (V^m(k) - \delta \ln \sigma^m(k)). \tag{41}$$

The value of moving can be substituted using (39) then (34) to yield

$$V^m(j) = \lambda[v_{zj} - \ln \sigma_{zj}], \tag{42}$$

which means the continuation value difference is represented with

$$V(j) - V(k) = (\lambda[v_{zj} - \ln \sigma_{zj}] - \delta \ln \sigma^m(j)) - (\lambda[v_{zk} - \ln \sigma_{zk}] - \delta \ln \sigma^m(k)). \tag{43}$$

The difference in value of moving to z from j vis-a-vis k is simply the difference in their moving costs from the origin

$$v_{zj} - v_{zk} = c_{zj} - c_{zk}. \tag{44}$$

This is the finite dependence step, in that the choice path can be returned to an equivalent point in the state space (here, a move to the location z). The successive substitutions have removed the actual value functions from the expression. Substituting in for v_{zj}, V_{zk} and collecting terms, we have

$$V(j) - V(k) = \delta(\ln \sigma^m(k) - \ln \sigma^m(j)) + \lambda(\ln \sigma_{zk} - \ln \sigma_{zj}) + \lambda(c_{zj} - c_{zk}). \tag{45}$$

Returning this to (40), we now have an estimating equation with only observed probabilities and parameters:

$$\begin{aligned} &\lambda[\ln \sigma_{jo} - \ln \sigma_{ko}] = \\ &(u_j - u_k) + (c_{jo} - c_{ko}) - \beta\delta(\ln \sigma'^m(j) - \ln \sigma'^m(k)) - \beta\lambda(\ln \sigma'_{zj} - \ln \sigma'_{zk}) + \beta\lambda(c_{zj} - c_{zk}). \end{aligned} \tag{46}$$

Moving versus staying

The approach for the move/stay decision is intuitively similar but requires additional substitutions and derivations to benefit from finite dependence.

The log odds ratio for staying versus moving is

$$\ln \sigma^s - \ln \sigma^m = \frac{1}{\delta}(V^s - V^m) \tag{47}$$

The difference in stay/move odds ratio between two different locations is then

$$\begin{aligned} &\delta[(\ln \sigma^s(j) - \ln \sigma^s(j)) - (\ln \sigma^s(z) - \ln \sigma^m(z))] = \\ &\qquad (V^s(j) - V^m(j)) - (V^s(z) - V^m(z)) = (V^s(j) - V^s(z)) - (V^m(j) - V^m(z)). \end{aligned}$$

Treat these separately as a moving block (the difference of moving values) and a staying block (the difference in staying values). The moving block is the difference

$$V^m(j) - V^m(z) = \lambda \ln[\sum_{k \neq j} exp(v_{kj})^{1/\lambda}] - \lambda \ln[\sum_{k \neq z} exp(v_{kz})^{1/\lambda}]. \tag{48}$$

By the same procedure as in (42), we can represent the moving value difference between j and z with respect to a third location k,

$$V^m(j) - V^m(z) = (v_{kj} - \ln \sigma_{kj}) - (v_{kz} - \ln \sigma_{kz}) = \lambda[(c_{kj} - \ln \sigma_{kj}) - (c_{kz} - \ln \sigma_{kz})]. \tag{49}$$

The staying block is the difference of flow utilities and continuation values,

$$V^s(j) - V^s(z) = u_j - u_z + \beta(V'(j) - V'(z)). \tag{50}$$

The continuation value difference returns us to (45), with z and k flipped:

$$V(j) - V(z) = \delta(\ln \sigma^m(z) - \ln \sigma^m(j)) + \lambda(\ln \sigma_{kz} - \ln \sigma_{kj}) + \lambda(c_{kj} - c_{kz}). \tag{51}$$

Putting the moving and staying blocks back together, we have

$$\begin{aligned}
&\delta[(\ln \sigma^s(j) - \ln \sigma^m(j)) - (\ln \sigma^s(z) - \ln \sigma^m(z))] = \\
&\qquad u_j - u_z + \beta\delta(\ln \sigma'^m(z) - \ln \sigma'^m(j)) + \\
&\qquad \beta\lambda(\ln \sigma'_{kz} - \ln \sigma'_{kj}) + \beta\lambda(c_{kj} + c_{kz}) - \lambda((c_{kj} - \ln \sigma_{kj}) - (c_{zj} - \ln \sigma_{zj})).
\end{aligned}$$

In general, the choice probability between one period and the next may differ as the state space evolves. Because our cross-sectional data uses only location as a state variables, we assume $\sigma_{kz} = \sigma'_{kz}$. This is acceptable when the variation in move probabilities between locations and types is more important than temporal variation, which we have argued elsewhere. Also, letting z be a normalizing location, we assume $c_{zj} = c_{zk}$. Then, collecting terms, we have

$$\begin{aligned}
&\delta[(\ln \sigma^s(j) - \ln \sigma^m(j)) - (\ln \sigma^s(z) - \ln \sigma^m(z))] = \\
&\qquad u_j - u_z + \beta\delta(\ln \sigma^m(z) - \ln \sigma^m(j)) + (\beta - 1)\lambda(\ln \sigma_{kz} - \ln \sigma_{kj} + (c_{kj} - c_{kz})).
\end{aligned} \tag{52}$$

Collecting the estimating equations

To collapse these into a system of equations, we collect the data terms for the lefthand side (target moments) and utility terms for the righthand side (model structure).

For the expression from the destination choice nest, (46), the equation divides as follows. The data target is

$$Y_1 = \lambda[\ln \sigma_{jo} - \ln \sigma_{ko}] + \beta\delta(\ln \sigma'^m(j) - \ln \sigma'^m(k)) + \beta\lambda(\ln \sigma'_{zj} - \ln \sigma'_{zk}). \tag{53}$$

The utility function is

$$u_j - u_k = \mu_j - \mu_k + \alpha R_j I(H = j) - \alpha R_k I(H = k), \tag{54}$$

which divides into the following structure and parameter matrices,

$$u(x_j) - u(x_k) = \begin{pmatrix} I(j) - I(k) & R_j I(H=j) - R_k I(H=k) \end{pmatrix}, \tag{55}$$

and

$$\theta_u = \begin{pmatrix} \mu_j - \mu_k \\ \alpha \end{pmatrix}.$$

The move cost term is

$$(c_{jo} - c_{ko}) = c(d_{jo}) - c(d_{ko}). \tag{56}$$

Notice the move cost intercept is differenced out and the distance term remains. Also, assuming $c_{zj} = c_{zk}$, the trailing move cost term falls out as well. The remaining distance term is the sum of the distance functions, dividing into the structure and parameter matrices as

$$\begin{aligned} d_{jo} - d_{ko} = & \\ & \begin{bmatrix} d_{ko}^{km} - d_{ko}^{km} & I(d_{jo}^{reg}) - I(d_{ko}^{reg}) & I(d_{jo}^{state}) - I(d_{ko}^{state}) & I(d_{jo}^{nhbor}) - I(d_{ko}^{nhbor}) \end{bmatrix}, \end{aligned}$$

and

$$\theta_c = \begin{pmatrix} c^{km} \\ c^{reg} \\ c^{state} \\ c^{nhbor} \end{pmatrix}. \tag{57}$$

For the expression from the move/stay nest, (52), the equation divides similarly, with different location indexing and scale parameters. The data target is

$$\begin{aligned} Y_2 = \delta[(\ln \sigma^s(j) - \ln \sigma^m(j)) - (\ln \sigma^s(z) - \ln \sigma^m(z))] & \\ - \beta\delta(\ln \sigma^m(z) - \ln \sigma^m(j)) - (\beta - 1)\lambda(\ln \sigma_{kz} - \ln \sigma_{kj}). & \end{aligned}$$

The utility function is

$$u(x_j) - u(x_z) = u_j - u_z = \mu_j - \mu_z + \alpha R_j I(H=j) - \alpha R_z I(H=z). \tag{58}$$

which divides into the following structure and parameter matrices,

$$u(x_j) - u(x_z) = \begin{pmatrix} I(j) - I(z) & R_j I(H=j) - R_z I(H=z) \end{pmatrix}. \tag{59}$$

and

$$\theta_u = \begin{pmatrix} \mu_j - \mu_z \\ \alpha \end{pmatrix}$$

The move cost term is

$$(c_{kj} - c_{kz}) = c(d_{kj}) - c(d_{kz}). \tag{60}$$

Notice the move cost intercept is differenced out and the distance term remains. The distance term is the sum of the distance functions, dividing into the structure and parameter matrices as

$$\begin{aligned} d_{kj} - d_{kz} = \\ &(\beta - 1)\lambda \begin{bmatrix} d_{kj}^{km} - d_{kz}^{km} & I(d_{kj}^{reg}) - I(d_{kz}^{reg}) & I(d_{kj}^{state}) - I(d_{kz}^{state}) & I(d_{kj}^{nhbor}) - I(d_{kz}^{nhbor}) \end{bmatrix} \end{aligned}$$

and

$$\theta_c = \begin{pmatrix} c^{km} \\ c^{reg} \\ c^{state} \\ c^{nhbor} \end{pmatrix} \tag{61}$$

Stacking these equations across locations and types yields the estimator represented by equation (26) and reported on in section 6.

D.4 Move Cost Intercepts

Move costs consist of an intercept term (for moving anywhere) and a distance term (depending on the origin-destination),

$$c_{kj} = c_0 + c(d_{kj}).$$

Notice that because the value of a location j from origin o is $v_{jo} = u_j + c_0 + c(d_{jk}) + \beta V(j)$, the value of moving term can be factored as

$$V_m = \lambda \ln[\sum_{i \neq o} exp(v_{io})^{1/\lambda}] = c_0 + \lambda \ln[\sum_{i \neq o} exp(u_i + c(d_{io}) + \beta(V(i))^{1/\lambda}]. \tag{62}$$

Hence, the odds ratio of moving versus staying can be represented as

$$\delta(\ln \sigma_s - \ln \sigma_m) = V_s - c_0 - \lambda \ln[\sum_{i \neq o} exp(u_i + c(d_{io}) + \beta(V(i))^{1/\lambda}]. \tag{63}$$

The moving cost intercept can be recovered from this odds ratio if one can account for the differences in the origin value $V_s, V_m - c_0$, which may make moving more or less likely in a given origin-type pair.

One way is to control generically for origins with fixed effects. That is, we can place the odds ratios on the lefthand side as target moments, and control non-parametrically for the differences across origin-type cells with a set of origin, type, or origin-by-type fixed effects. Another way is to derive the value using the structure of the model, accounting for the origin value. Notice that the odds ratio is

$$exp(V^m(o))/exp(V^s(o)) = \frac{(\sum_{i \neq o} exp(v_{io})^{1/\lambda})^\lambda}{(exp(V^s(o))^{1/\lambda})^\lambda}, \tag{64}$$

which is expanded as

$$\frac{(\sum_{i \neq o} exp(v_{io})^{1/\lambda})^\lambda}{(exp(V_s)^{1/\lambda})^\lambda} = (\sum_i exp(v_{io} - V^s(o))^{1/\lambda})^\lambda = (\sum_i exp(V(i) + c_o + c(d_{io}) - V^s(o))^{1/\lambda})^\lambda. \tag{65}$$

Notice that the c_0 term, being the common intercept across all locations, can be factored out as before. Then, because $v_{iz} = V(i) + c_{iz}$, the difference in location values can be written $V(i) - V(o) = (v_{iz} - c_{iz}) - (v_{oz} - c_{oz})$, and we can write the above as

$$\frac{(\sum_{i \neq o} exp(v_{io})^{1/\lambda})^\lambda}{(exp(V_s)^{1/\lambda})^\lambda} = exp(c_0)(\sum_i exp((v_{iz} - v_{oz}) - (c_{iz} - c_{oz}) + c(d_{io}))^{1/\lambda})^\lambda.$$

The difference in values between i and o is their odds ratios in the choice probability originating from z, $exp(v_{iz} - v_{oz}) = \frac{\sigma_{iz}}{\sigma_{oz}}$, which means this can be written as

$$\frac{(\sum_{i \neq o} exp(v_{io})^{1/\lambda})^\lambda}{(exp(V_s)^{1/\lambda})^\lambda} = exp(c_0)(\sum_i \frac{\sigma_{iz}}{\sigma_{oz}} exp(-(c_{iz} - c_{oz}) + c(d_{io}))^{1/\lambda})^\lambda. \tag{66}$$

Hence, we have the log odds ratio estimating equation,

$$\ln \sigma^m - \ln \sigma^s = c_0 + \lambda \ln(\sum_i \frac{\sigma_{iz}}{\sigma_{oz}} exp(-(c_{iz} - c_{oz}) + c(d_{io}))^{1/\lambda}). \tag{67}$$

The choice probabilities are observable from the data, and the distance-dependent cost terms were recovered in the utility parameters estimation step, meaning the only remaining unknown term is c_0. Each origin location and birthplace-type cell provides an observation of the move/stay odds ratio. Constructing the term within the summation for each origin o accounts for the differential value the cell type encounters when facing a system of locations, just as a control function would account for unspecified error terms. The residual average from the comparison of move/stay odds ratios provides the estimate of the move cost intercept.

The parametric and nonparametric approaches produce similar results. We prefer the parametric control approach because it maintains consistency with the model throughout the analysis.

D.5 Scale Parameters

There are scale parameters for each level of nest, δ for the move/stay decision and λ for the destination (conditional on moving) decision. The relationship between them is described by equation (14), comparing inflow rates to future move out rates:

$$\begin{aligned}\lambda[\ln \sigma_{jo} - \ln \sigma_{ko}] + \beta\lambda(\ln \sigma'_{zj} - \ln \sigma'_{zk}) = \\ (u_j - u_k) + (c_{jo} - c_{ko}) + \beta\delta(\ln \sigma'^{m}(k) - \ln \sigma'^{m}(j)) + \beta\lambda(c_{zj} - c_{zk}).\end{aligned}$$

Because the utility is only identified up to scale, we cannot separately identify these parameters, but this comparison elicits their relative importance. We set $\lambda = 1$ and then seek to estimate δ to set the scale ratio. By assumption, $c_{zj} = c_{zk}$. Then, controlling non-parametrically for the two potential destinations j, k and for their pairing with origin o – i.e, running a gravity-style regression with destination and origin/destination fixed effects – we can recover the δ parameter as the coefficient on the difference in move out rates.

$$\frac{1}{\beta}[\ln \sigma_{jo} - \ln \sigma_{ko}] + \lambda(\ln \sigma'_{zj} - \ln \sigma'_{zk}) = \frac{1}{\beta}(u_j - u_k) + \frac{1}{\beta}(c_{jo} - c_{ko}) + \delta(\ln \sigma'^{m}(k) - \ln \sigma'^{m}(j)). \quad (68)$$

The regression is run for each demographic type τ in the model. In practice, we do the scale parameter term as the very first step in estimation, because we need the δ parameter to set up the utility parameter estimation step.

D.6 Data Smoothing

Even with a relatively large dataset and large LLMs, migration is infrequent enough that a fully interacted cell definition resulted in many empty cells. Granularity is a challenge here. Once the data are cut to, for example, 40-something college-educated workers living in Houston but born in Cleveland, there are few individuals populating the cell. We may fail to observe any of this type moving to, say, Kansas City, but do not believe that the probability of that event is literally zero. Our smoothing procedure is designed to make aggregated cells that preserve the kinds of detail in the stylized facts presented above.

The smoothing procedure is to first run the raw data through a Poisson regression of flows on a rich set of destination by type and origin by type fixed effects, as well as a distance function. Under the Poisson regression, the expected number of flows is modeled as

$$\sigma_{jo}^{\tau,H} = \exp[A_j^{\tau,H} + B_o^{\tau,H} + C^{\tau,H}(d_{jo})],$$

where A_j are destination fixed effects and B are origin fixed effects, each of which is indexed by type. C is a cost function to account for resistance between markets imposed by distance. In practice, we use age (in decades) by education (college/noncollege) by home status (at home, not at home, foreign born, and from the outside option location). This is somewhat coarser than a full suite of all birthplace-type fixed effects, but that is because we need some observations in the cell to be able to impute a flow probability, which is not satisfied in all origins for some small birthplaces. The cost function is the four-dimensional specification we use in our main utility estimation, interacted with educations, and for the imputation procedure only, with a term interacting distance in kilometers with home status (i.e., if the destination is home).

We use a Poisson Pseudo-Maximum Likelihood (PPML) estimator contributed by Correia et al. (2020) because it can accommodate high dimensions of fixed effects. We use the projections from the PPML regression to impute flows for each cell, and then calculate the choice probabilities using the imputed flows.

The advantages of the Poisson regression are twofold. First, it can handle zeroes in the raw data and produce a smoothed, nonzero estimate of the number of flows – and unlike actual migration data, it can produce non-integer estimates of flows, including flows above zero but below one. Second, it is a parametric smoothing procedure that is in practice robust to context-specific researcher choices such as the number of locations and types to include. We also tried a nonparametric smoothing procedure that aggregated cells, calculating marginal probabilities, and then produced conditional probabilities as imputed outputs. This produced imputations highly correlated to those from the Poisson procedure, but the latter method is less dependent on elective actions, such as how to aggregate the cells.

Figure E1: Simulated Migration Rate Over Time, by LLM Group

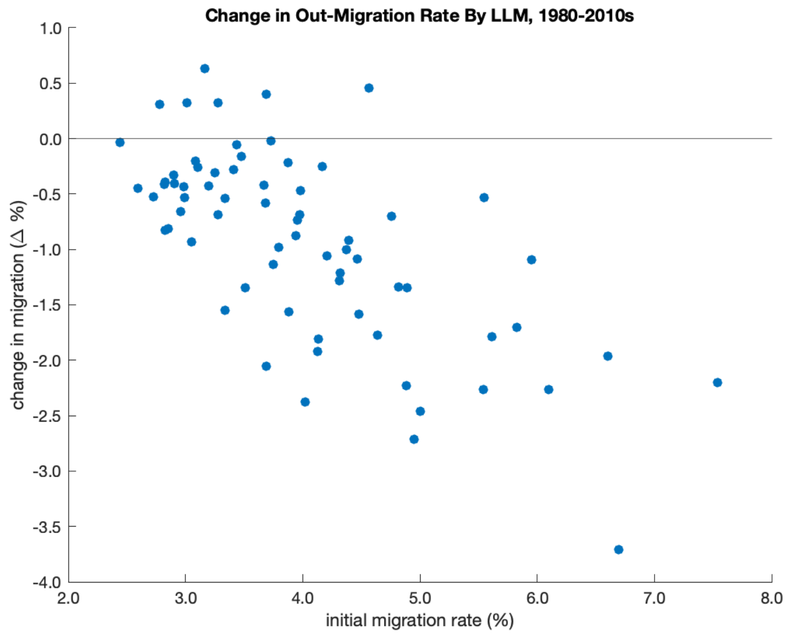

NOTES: The figure plots the change in migration in the simulated model, 1980 to 2010s average, against the baseline simulated rate (2010s), using the preferred specification (4). (Source: Authors' calculations using simulation results.)

E Additional Model Simulation Results

This section presents additional results from the model simulations covered in Section 8.

Figure E1 plots the change in change in migration from 1980 to the 2010s average under the simulation. The model is able to generate a negative correlation between average mobility and the decline in mobility while allowing for a fair degree of locational heterogeneity.

Figure E2 shows the model's simulated migration rates by demographic group (education by age categories). It shows that each group follows a hump-shaped pattern, with a peak in migration from 1960 to 1990, depending on the group. The magnitude and timing of the decline varies by group, but each shows a decline in recent decades. The ubiquity of the decline across groups was part of the puzzle presented by earlier papers on the migration decline (such as Molloy et al. (2011)).

Figure E3 shows the timepath of the subtotal simulations broken out by LLM category. All categories show peak migration around 1980. For medium and slow locations, this is virtually all due to demographics (aging). Fast locations were affected by a rise in home attachment, which shows a clear hump in migration along with the demographic trends. These two effects combine to generate a substantial decline in migration from 1980 to the 2010s.

Figure E2: Simulated Migration Rate Over Time, by Demographic Group

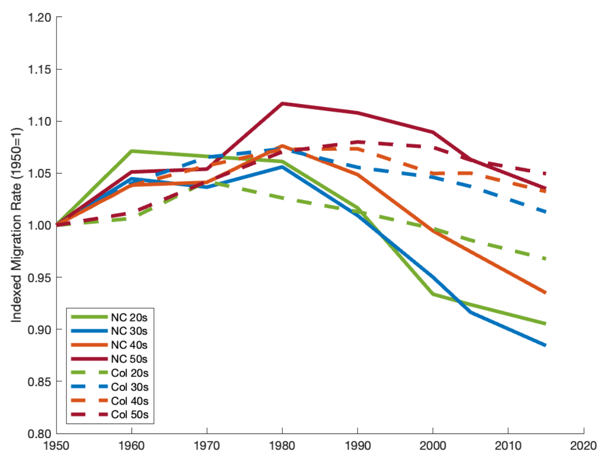

NOTES: The figure plots the simulated timepath of migration rates by demographic group using the preferred specification (4). (Source: Authors' calculations using simulation results.)

Figure E3: Decomposition of Simulated Migration Rate Over Time, by LLM Group

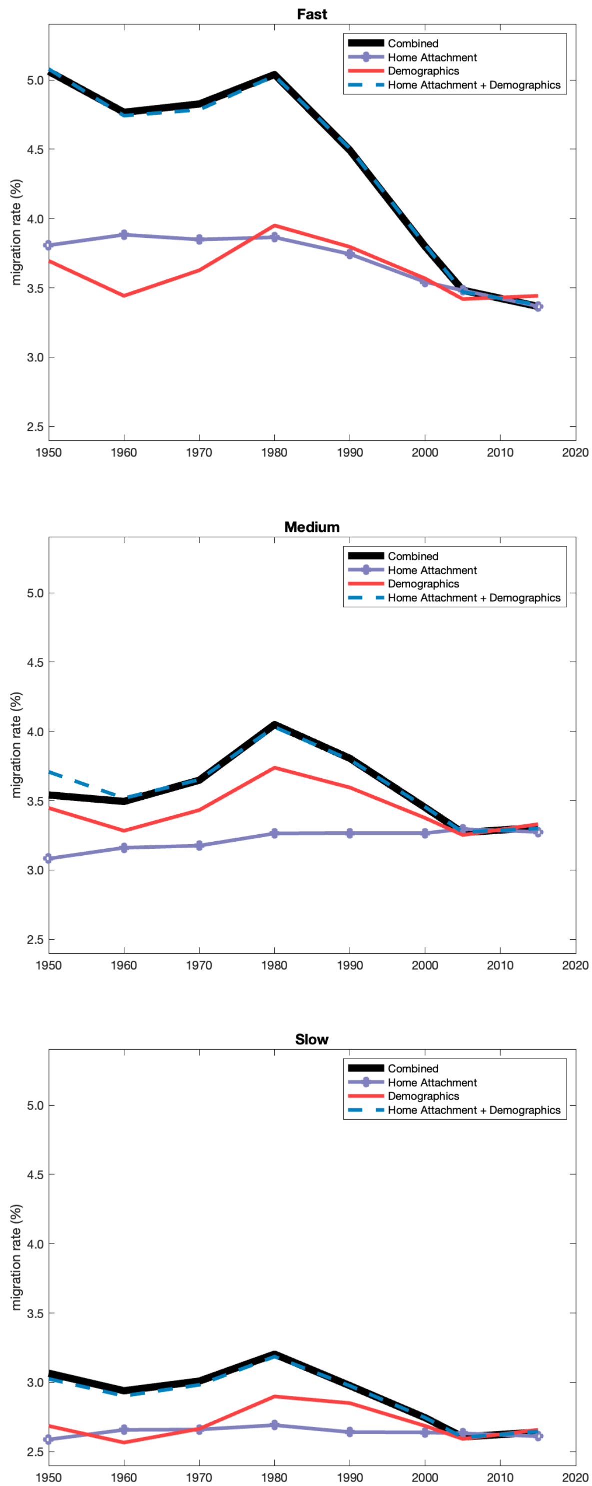

NOTES: The figure plots the simulated timepath of migration rates for counterfactual simulations as denoted by the legend. Each panel corresponds to a different LLM category. All figures use the same vertical scale. The differences from 1980 to 2010s average are reported in Table 6. (Source: Authors' calculations using simulation results.)