House Price Contagion and U.S. City Migration Networks

Gregor Schubert
Harvard University *
April 26, 2021

Job Market Paper
(Click here for most recent version)

Abstract

Why do national trends in house prices spread more to some cities than to others? This paper proposes an explanation of house price contagion based on migration spillovers between U.S. cities: Increases in house prices as a result of local economic shocks and housing supply constraints drive out-migration to other cities. These migration flows are more likely to affect cities with stronger pre-existing migration links to the origin cities, and increase house prices at these destinations. I use the network structure of inter-city migration to develop an instrument for identifying causal spillover effects between cities: I find that an increase in other cities' house prices by 10% in the long run causes a 6.3% house price move in a city exposed to the shock through migration links. I show that migration spillovers from the effect of interest rate declines on house prices in other cities can explain 32% of the cross-sectional variation in house price growth during the run-up to the housing boom of the 2000s. To quantify the effect of changes in migration costs and housing supply constraints on these house price spillovers, I develop and estimate a dynamic spatial equilibrium model that incorporates forward-looking migration choices. After estimating this model with U.S. data, I show that lower migration costs substantially reduce the dispersion in house price growth: without worker mobility, the spread in house price growth across cities in response to wage shocks would be 65-70% larger. Moreover, declines in worker mobility increase the impact of housing policy on the distribution of house price growth across cities.

*Email: gschubert@g.harvard.edu. Website: `https://scholar.harvard.edu/gregorschubert`. I would like to thank my advisors Ed Glaeser, Larry Katz, Jeremy Stein and Adi Sunderam for their guidance and support. This paper also benefited immensely from discussions with Pol Antras, Adrien Bilal, John Campbell, Gabriel Chodorow-Reich, Marco Di Maggio, Fabian Eckert, Emmanuel Farhi, Xavier Gabaix, Caitlin Gorback, Robin Greenwood, Sam Hanson, Myrto Kalouptsidi, Robin Lee, Marc Melitz. I would also like to thank participants of the Harvard Labor Lunch, Industrial Organization Workshop, International Trade Workshop, Macro Workshop, and Finance Workshop, as well as participants of the UEA Ph.D. Workshop 2020 for many helpful comments and suggestions. I gratefully acknowledge receiving financial support from the Bradley Foundation, Harvard Business School, and the John R. Meyer Dissertation Fellowship of the Harvard Joint Center for Housing Studies while working on this research.

1 Introduction

House price growth during national boom and bust cycles varies greatly across cities in the U.S. During the 2000-2007 housing boom, real house prices in Boise, ID, grew by 41%, while house price growth in similarly sized Midwestern cities like Wichita, KS, or Jackson, MS, was only 2% and 9%. Such large cross-sectional differences are not unusual: the spread in 2000-2007 real house price growth between 0% at the 10th percentile and 91% at the 90th percentile is far larger than the national average of 40%[1] that attracted substantial policy attention. Even after accounting for the direct effect of land constraints and their interaction with national macroeconomic trends, these geographic differences are substantial as Panel (a) of Figure 1 shows. Spatial gaps in cities' co-movement with national house price trends also persist over longer time periods. The map in Panel (b) of Figure 1 shows the local correlation with the national cycle (the house price "beta") over the 1990-2017 period, which ranges from 0.4 at the 10th percentile to 1.9 at the 90th. Why do national cycles in house prices affect some U.S. cities more than others? How do shocks to local economies spread between cities? And why did some cities experience high house price growth during the 2000s and 2010s even in the absence of large local economic shocks?

This paper aims to answer these questions by providing evidence for the importance of *migration* as an important causal channel for spillovers between cities. In particular, I make two contributions: First, I show new empirical evidence that house price growth in one city can have a causal spillover effect on other cities through migration flows, which partly explains house price growth in cities that don't experience direct economic shocks. Second, I construct and estimate a model of location choices with forward-looking agents that allows me to quantify the role that mobility plays in mitigating extreme house price growth outcomes across cities.

The proposed mechanism for why house price growth spills over to some cities more than to others is the following: First, economic shocks benefit one group of workers in a city and drive up local house prices if housing supply is constrained. This causes other workers to leave the city in pursuit of more affordable housing elsewhere. These outflows are more likely to affect cities with strong migration links to the origin city – and the increased population inflows drive up those destinations' housing demand and prices. As a result of such migration spillovers, high house price growth may spread even to cities that do not receive fundamental shocks themselves, but are highly exposed to outflows from cities that do. Returning to the example from above, the high house price growth in Boise, ID, can in part be explained by the fact that its population grew by 13% over 2000-2007, while that of Wichita, KS, and Jackson, MS, stagnated. One reason for this is that Boise is an important destination for people leaving expensive booming cities in California, such as Los Angeles, San Francisco, or San Jose. These three metro areas alone were responsible for 2000-2007 net migration into Boise corresponding to almost 4% of its population.[2]

In the first part of the paper, I start by establishing stylized facts about the networks of migration between U.S. cities: I show that, for most cities, only a small number of other cities are destinations or origins for substantial numbers of migrant outflows or inflows. Moreover, this set of relevant other cities does not change much over time. This persistence can be explained by migration costs that are a function of enduring city characteristics. As a case study of how

[1] Averages and percentiles are computed across population-weighted commuting zones.

[2] For comparison, in the 1990-1999 period that preceded most of the housing boom, Boise, ID, was among the Top 50 migration destinations for the Los Angeles, Oakland-San Francisco, San Jose, and Santa Barbara metro areas. In contrast, during the same time period, Wichita's and Jackson's highest destination rankings for *any* Gyourko et al. (2013) "superstar" city were 81st and 173rd, respectively, and they saw 2000-2007 population growth of -2% and 1%, respectively.

Figure 1: House price growth and beta. Panel (a) shows a map of commuting zone average (cumulative) real house price growth over 2000-2007, residualized with regard to the effect of annual time trends, local land constraints, and the interaction between them in a population-weighted regression. Panel (b) shows house price growth "betas" for the annual house price growth effect in each period, computed as the slope coefficient in a regression of each CZ's annual house price growth series 1990-2017 from Federal Housing Finance Authority data on the series of leave-one-out average growth in each year.

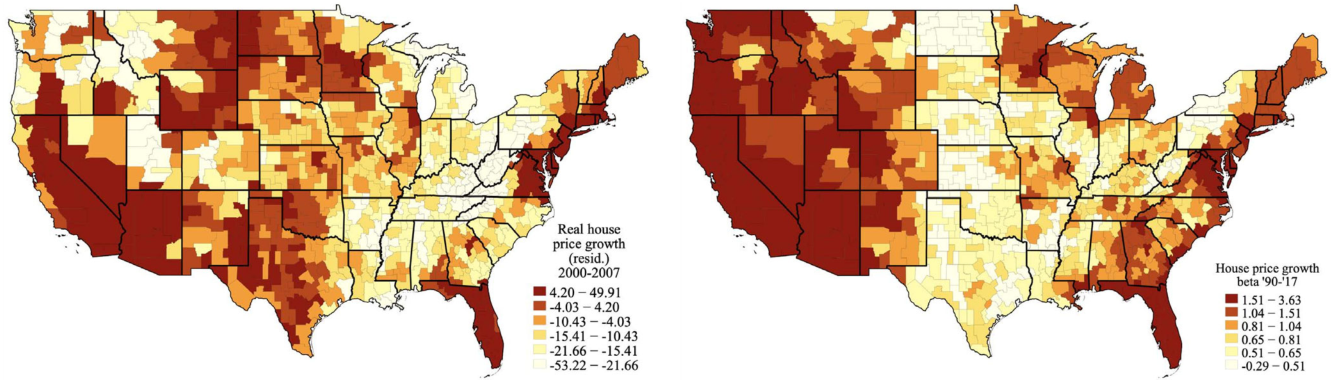

(a) Resid. real house price growth 2000-2007 **(b)** House price growth beta 1990-2017

migration spillovers between cities might operate, I demonstrate that the ex ante defined top destinations of migrants from booming "superstar" cities (Gyourko et al., 2013) were likely to experience particularly large housing booms and high population growth during the housing cycle of the early 2000s.

Then, I develop an empirical approach to identify causal reduced-form spillover effects between cities. I use the network structure of migration flows to measure each city's migration exposure to other cities' housing markets and estimate spillover effects from the degree to which a focal city's house price growth co-moves with house price shocks in other migration-linked cities. To identify exogenous variation in house prices in a city's migration network, I construct an instrument consisting of Bartik (1991) type wage shift-share shocks interacting with land constraints in *other* cities – and exposure to those shocks through the same migration network links. In the estimation I control for any *direct* effects of those shift-share shocks on each focal city to ensure that any house price effects come only from indirect spillovers via migration links.

Using this "network IV" approach, I find that causal network spillover effects are sizeable. In the long run, a 10% increase in house prices in all other migration-connected cities results in a 6.3% increase in house prices and a 1.9% increase in the population of a city affected through the migration spillover channel. The effects on population and house prices increase for two years after the initial shock before levelling off at these long-run values. These spillover effects on house prices are stronger in cities that are more land-constrained: an additional 10 ppt of local land unavailable for construction increases the spillover effect from a 10% house price shock by 0.9 ppt.

In addition, I estimate the spillover effects on mortgage credit and construction – which support the proposed migration mechanism. Mortgage loan originations are significantly higher for several years in response to a migration spillover shock. Furthermore, the house price response to migration spillovers can partly be explained by congestion in the construction sector: construction permit volume responds to migration spillovers only with a delay and remains 31% above its initial level even after 5 years in response to a 10% house price shock in other cities. As additional evidence of construction lags, I show that housing vacancies decline in the cities most affected by housing booms in the 2000s and 2010s while regional time-to-build delays in housing construction increase procyclically.

How important are these migration spillover effects? I quantify the spillovers between cities that arise from two different shocks that the literature has shown to have had direct effects on house prices in the run-up to the housing boom of the 2000s:[3] (1) Declining real interest rates, and (2) industry wage shocks. I show that these shocks had indirect spillover effects on cities that have strong migration links to cities that received large direct shocks. Spillovers from interest rate shocks affecting other cities can explain 32% of the cross-sectional variation in actual house price growth, while spillovers from wage shocks to other cities explain 12%. I also provide new evidence that migration links individually outperform in predicting high inter-city house price correlations when compared to measures of physical distance, social network links, trade flows, or similarity in industry structure between cities.

The second part of the paper focuses on how the migration spillovers between linked cities affect the aggregate cross-sectional pattern of house price growth in the U.S. While the reduced-form estimates take structural parameters as given, the recent decline in U.S. domestic migration rates and increase in housing supply constraints raise the question of how changes in mobility or house price elasticities affect house price dynamics in the presence of migration spillovers.[4] To study these counterfactual questions, I develop and estimate a dynamic spatial equilibrium model that allows me to quantify the role of migration costs and housing supply constraints in generating the observed distribution of house price growth. The model allows for heterogeneity in housing supply elasticities across cities, forward-looking location choices by migrants, and heterogeneity in location preferences between education groups. By capturing the degree to which cities are interconnected, this model can account for rich geographic differences in spillovers of economic shocks. I use instruments based on wage shift-share shocks to identify the key parameters characterizing location choices, industry choices, labor demand and heterogeneity in housing supply constraints from the structural equations of this model.

An important estimation challenge is that moving costs mean that workers cannot re-optimize their location costlessly as circumstances change. As a result, moving decisions can depend on the unobserved expected option value of being in a particular city at a particular point in time. Not accounting for these expectations about future opportunities in different cities could lead to omitted variable bias in estimating the sensitivity of location choices to observable characteristics. In estimating the location choice parameters, I therefore use a "renewal state" approach, which differences out the future option values of locations: I assume that the future utility of being in a particular location is independent of the migration path taken to get there. As a result, comparing characteristics between two paths that end up in the same location holds the option values at the end point constant, which I exploit in the estimation of the location choice sensitivity parameters.

To estimate this model in a recent U.S. context, I have to overcome the challenge that migration flow data by education group is noisy and only available for 2005-2017. I reduce noise in the migration data by using an empirical Bayes shrinkage estimator to combine raw survey data from the American Community Survey with a LASSO prediction of smoothed migration flows. I also use the model structure together with observable aggregate data to impute historical flows by

[3] See, e.g. Saks (2008) or Saiz (2010) for examples showing direct effects of industry shift-share shocks on local house prices, and Himmelberg et al. (2005) or Glaeser et al. (2012) for studies showing a direct effect of interest rate changes in conjunction with land constraints on local house prices.

[4] See Molloy and Smith (2019) for a summary of the evidence showing a secular decline in migration, and Ganong and Shoag (2017) or Gyourko et al. (2019) for measures that show increases in housing supply constraints.

education group for the year 2000.

Last, I use the model to conduct counterfactuals for the effect of changes in migration costs and housing supply constraints. I quantify the counterfactual impact of wage shocks on the distribution of house price growth across cities during the 2000-2007 and 2012-2017 housing boom periods, under different assumptions about migration costs or housing supply constraints. I find that migration spillovers play an important role in spatially distributing the house price effects of concentrated economic shocks. My counterfactual simulations show that less mobility results in more dispersed house price growth: with prohibitive migration costs, – no mobility – measures of the spread in house price growth across cities in response to wage shocks would be 65-70% larger than under observed baseline migration costs. That is, fewer cities experience extremely high (or low) house price growth after a positive shock if workers can move more easily between cities. I show that this effect on the dispersion of of house price growth is qualitatively similar to that of a reduction of all above-median housing supply constraints to the median. Mobility also affects the distribution of house price effect betas: higher costs of moving between cities lead to heavier tails of cities with house price growth that is highly positively or negatively correlated with the national trend.

Differences in how mobile workers are across cities also interact with the effectiveness of housing policy: decreases in mobility almost universally *increase* the degree to which reductions in supply constraints decrease the dispersion of house price growth. If the secular decline in observed U.S. inter-city mobility over the last three decades[5] represents a fundamental shift in the ability of workers to move, this analysis implies that policy changes to reduce housing supply constraints would be *more* effective today at reducing the dispersion in house price outcomes than in the past. Conversely, policies that enable greater worker mobility through lower migration costs can mitigate the impact of tighter housing supply constraints that widen the gap in house price growth across cities.

These findings have important implications for policy-makers and stakeholders in real estate markets: In the aggregate, declines in workers' ability to move across cities are predicted to lead to greater cross-sectional differences in house price growth between cities, including more housing markets with relatively large house price increases. At the same time, my analysis suggests for individual housing markets that migration links to other cities that are experiencing economic booms help to predict increases in local housing demand; and being able to better predict local surges in demand can benefit both real estate investors and urban planners trying to predict housing needs and zoning requirements. Moreover, to the degree that the migration spillover mechanism explains part of the volatility in house prices in U.S. cities like Phoenix, AZ, or Boise, ID, it also informs financial regulators and macroeconomic policy-makers evaluating to what degree local asset price changes are the result of speculation or credit supply changes rather than driven by demand fundamentals.[6]

Related literature. This work builds on a number of papers that have documented housing market connections between cities: DeFusco et al. (2015) show that contagion appears to have played a role in the expansion of the 2000s boom across local housing markets. In contrast to this paper, they focus on closest neighbors and do not consider migration as an important spillover channel in their analysis. My analysis also complements Chinco and Mayer (2016), who show that out-of-town speculators, especially coming from large cities, play an important role in driving

[5]See Panel (a) of Appendix Figure A1 for a plot of the 1990-2017 trend in gross migration rates.

[6]See, e.g. the "Financial Stability Report" by Board of Governors of the Federal Reserve System (2020) for an example of national policy-makers tracking house price increases in local markets.

house price appreciation in the mid-2000s. If spillover migrants are accompanied by out-of-town speculators from large cities, my paper provides one explanation for why some cities are more likely to be the target of out-of-town speculators than others. Moreover, the *cross-city* migration mechanism that I am studying is similar to the *within-city* gentrification dynamics shown by Guerrieri et al. (2013). They show that neighborhoods that are closer to rich neighborhoods are more likely to experience house price increases after positive economic shocks to a city, providing implicit evidence of within-city migration as the driving force of this dynamic. In contrast, I study the spillover effect of positive shocks *across* proximate cities, where "proximity" is defined by the ease of migration. Bailey et al. (2018a) also study network connections between cities – in the form of social media friendships – and their effect on individual house purchases. However, they focus on heterogeneity in online social networks among different residents *within* one county. In this paper, I try to explain overall house price dynamics for a large network of cities using a mechanism explicitly focusing on the effect of movers, which may complement spillovers through social media connections. Cotter et al. (2015) document large correlations in housing returns between cities, and Sinai and Souleles (2013) showed that there are positive house price correlations across migration-linked cities. I provide a mechanism for how migration links can cause house price correlations, as well as new evidence that compares migration links to other measures of inter-city links.

This paper is also closely tied to a large literature on differences in house price growth and mortgage credit origination across geographic areas.[7] For example, a rapid expansion in mortgage credit in some areas that drove house price increases during the run-up to the housing boom of the 2000s has been tied to the presence of subprime borrowers (Mian and Sufi, 2009), bank deregulation (Favara and Imbs, 2015), credit supply subsidies from government-sponsored enterprises in conjunction with differences in financial integration (Loutskina and Strahan, 2015), and changes in conforming-loan limits (Adelino et al., 2012). Moreover, regional differences in house price growth have been shown to arise from the interaction of national changes in real interest rates with local housing supply constraints (Himmelberg et al., 2005; Glaeser et al., 2012), and a concentration of high-income households in supply-constrained cities (Gyourko et al., 2013). This paper shows that all of these shocks can have indirect effects through migration spillovers on other cities that are not directly exposed to them. At the same time, I propose differences in migration exposure to house price and credit shocks in other cities as a new explanation for the geographic variation in the expansion of mortgage credit and house price growth.

In this paper, I also add to a literature that links migration flows and house price dynamics. For example, Howard (2020) shows that migration flows into a city stimulate the local economy and house prices, and Boustan (2010) shows that the arrival of Blacks during the Great Migration had a negative effect on city housing demand by whites. While this literature focuses on the effect on the city receiving migrants, I explicitly connect house price dynamics in origin and destination cities and quantify the aggregate effects of such spillovers during historical boom periods. Moreover, my model allows me to quantify the role of changes in structural parameters, such as migration costs, in changing house price dynamics, which is not possible in reduced-form approaches.

I also contribute to a literature on the methodology of spatial equilibrium and dynamic discrete choice models: my dynamic model builds on forward-looking location choice models like those used by Diamond et al. (2017) and Almagro and Domínguez-Iino (2020), as well as the renewal action approach for estimating location choice parameters as applied by Scott (2013), and generalized

[7]There is also a voluminous literature trying to explain variations in the housing cycle within cities (Mian and Sufi, 2009), or the size of the national cycle (e.g Landvoigt (2017); Kaplan et al. (2020)) which does not explicitly focus on the cross-sectional geography of differences in the housing cycle across cities.

by Kalouptsidi et al. (2020). My method for inferring stationary equilibria in spatial equilibrium models with migration costs combines methods developed by Ahlfeldt et al. (2020) and the "dynamic hat" simulation approach in Caliendo et al. (2019), to obtain transition paths to steady states for dynamic forward-looking spatial equilibrium models. While the intuition that shocks can "spill over" between cities is common to many papers with spatial equilibrium dynamics and location choice, my contribution consists of both quantifying the size of these spillover effects through a migration channel, and applying this intuition to understanding historical patterns of house price contagion and cross-sectional variation in housing booms, where migration patterns have previously not been considered as a driving force.

2 U.S. City Migration Networks: Descriptive Evidence

This section shows that there are meaningful differences in migration networks between U.S. cities, what drives these differences, and that they are persistent over time. Moreover, I show that identifying cities ex ante that are likely destinations for migration flows from cities that are likely to receive large positive economic shocks is predictive of high population and house price growth during the housing boom of the 2000s.

2.1 Migration networks between U.S. cities are sparse and persistent

When I am talking about the "U.S. city migration network", I am referring to the size of migration links between commuting zone ("city") nodes. These links are characterized by the number of migrants flowing from an origin city to a destination city. Moreover, we can distinguish between inflow and outflow networks, which can be very different for the same reference city.

To illustrate the geographic variation in migration networks, Figure 2 shows the migration network for the 2000-2007 housing boom period for inflows into Boston (left panel) and Dallas (right panel). For comparison, all inflows have been scaled to be shares of all inflows into the city from the continental U.S., with darker colors indicating a greater share of inflows coming from that origin commuting zone. Comparing the two maps, four things are immediately obvious: (1) Larger population centers naturally have a greater migration impact on other cities. (2) Migration networks can differ substantially between cities. (3) Migration flows are not limited to within-region flows, as many strong links extend to the farthest corners of the U.S. (4) Migration networks are sparse – that is, a majority of possible CZ-to-CZ links are too small to be recorded in the IRS data, while large flows concentrate into a small number of origin CZs for each destination. To show that these patterns are not unique to Boston and Dallas, Appendix Figure A2 shows the migration networks for a number of other large CZs.

Number of migration links. We can also document these connectivity patterns more systematically by computing the average number of inflow links that cities have, where a "link" is defined as an annual migrant inflow from another city recorded by the IRS.[8] These network characteristics are shown in Table 1 for the 1990-2010 period.[9] In the first row, I document that CZs on average receive migrants from 30 other CZs per year from a potential total of 722 CZs in the

[8] The threshold for the IRS to record the name of origin cities are flows corresponding to at least 10 tax returns.

[9] In 2011, the IRS changed its methodology for computing gross migration rates, leading to noise in the data, which is why I do not include post-2011 here. However, net migration rates do not seem to be affected by the methodology change, which is why I use data through 2017 in the regressions.

continental U.S.,[10] which shows that migration networks are sparse. That is, sizeable migration only occurs between particular city pairs rather than being widely dispersed.

Persistence. These migration links are also highly persistent over time – indicating that they are not driven by temporary shocks: even at a 10-year horizon, on average 80% of the origin CZs of incoming flows remain the same. Moreover, persistent migration networks do not only exist at a local level: the second and third rows of Table 1 show that, in an average year, cities receive inflows from 26 CZs that are at least 50 miles away, and even at 150 miles distance, the average city still has 21 links, of which 62% persist 10 years later. Comparing median and mean links, we can see that migration connectivity with other cities is skewed: many cities only have a handful of long-distance migration links, while a few (usually large) cities are widely connected.

Characteristics of movers. In order to understand which population groups are driving U.S. migration patterns in general, I consider evidence from the Current Population Survey in Appendix H.3 on the demographics of inter-county movers in the U.S. I find that inter-county migration is driven mainly by young workers, and by workers who are employed. In contrast, older age groups' share of migration is less than their population share and retirees make up less than 10% of intercounty movers.

2.2 Migration links are predicted by city similarity

Why are migration networks persistent? One explanation may be that migrants belonging to a particular demographic group are attracted to cities where other members of their group have gone before. There is a wealth of evidence that this is the case for *international* migration. For instance, Bartel (1989) shows that immigrants to the U.S. tend to cluster in cities that already have a high share of immigrants of the same ethnicity.[11] A similar mechanism of path dependence has been found for historical domestic migration patterns (see, e.g. Boustan (2010)).[12] A possible explanation for persistent migration patterns may be time-invariant city characteristics that shape both historical and later moving patterns in a similar way. For instance, Fishback et al. (2006) show that natural amenities played a role in U.S. domestic migration in the 1930s.[13]

As I am not aware of a recent study that systematically explores the correlates of migration costs over the last two decades, I analyse whether this is the case in the U.S. in this section.

Migration gravity model. There are two separate descriptive questions to answer: (1) What factors correlate with a city's attractiveness to migrants from *any* location? (2) Holding city attractiveness constant, *which* cities are more likely to be part of the same migration network than others?

[10]Based on 1990 CZ definitions.

[11]In fact, this insight has seeded a large empirical literature in economics which uses historical immigration flows between geographies as predictors of later flows (e.g. Card (2001)). One proposed explanation for why immigrants follow in each other's footsteps is that group connections to previous migrants decrease uncertainty about the destination and lower moving and adjustment costs (Carrington et al., 1996).

[12]Boustan (2010) uses county economic characteristics in 1940-1970 to predict out-migration from the U.S. South and combines this with historical migration patterns to predict inflows into Northern cities. She explains the persistence in migrant settlement patterns as being at least partially due to fixed train routes and community networks.

[13]For the more recent period of the 1990s, Glaeser and Shapiro (2001) find that cities with warmer and drier weather were able to attract more migrants, and Glaeser et al. (2001) show that, in addition to weather, U.S. city growth since the 1970s also correlates with consumption amenities, such as coastal proximity, restaurants, and performance venues.

To answers these questions, I first estimate a simple gravity model of migration of the form

$$\ln Mig_t^{i\to j} = \alpha_t + \underbrace{\gamma_1 \ln dist_{ij} + \beta' \|\mathbf{X_j} - \mathbf{X_i}\|}_{\text{Migration cost}} + \underbrace{\theta_{it} + \theta_{jt}}_{\text{City attractiveness}} + \epsilon_{ijt} \tag{1}$$

where I decompose log annual migration flows between each city pair into a national trend α_t, various bilateral distance and difference parameters that capture the cost of migrating between cities; and fixed effects that capture the overall attractiveness of a city, i.e. its overall tendency to receive or send migrants. I estimate this equation using Poisson Pseudo-Maximum Likelihood on the total sample of 1990-2017 IRS migration flows between CZs. To explore differences by education group, I also estimate the migration cost parameters on a shorter 2005-2017 panel of flows by college and non-college workers, which is imputed from ACS and IRS data as described in Appendix F.6.

Migration cost determinants. Migration costs are parameterized to depend on various bilateral components. The most basic proxy for geographical distance is log physical distance between origin and destination. In addition, I allow migration costs to depend on cities being in the same state or Census region, as well as interactions of these institutional boundaries and physical distance. Informed by the migration literature discussed above, I also include absolute differences between cities in the population share of non-mainline Christian denomination[14] and ethnicity shares (as proxies for culture), water access and mean January temperatures (climate and natural amenities), and 2-digit industry shares (economic structure).[15]

The estimated coefficients for the migration cost determinants are shown in Table 2, with negative coefficients indicating that greater distance or absolute difference imposes a greater cost on migrants. Overall, almost all the included city distances and differences have a significant negative effect on the ease of bilateral migration.[16] This rationalizes persistent differences in migration networks between cities: for a given level of city attractiveness, migrants are more likely to flow to and from cities that are more similar, or closer to one another.[17]

Heterogeneity in migration costs. The results by education group (see columns 2 and 3) highlight that the aggregate migration patterns obscure differences between groups. I find that physical distance and state / regional boundaries matter more for non-college workers.[18] Conversely, college workers seem to put a greater weight on similar industry composition and the prevalence of non-traditional religious denominations. In Appendix Section H.1, I additionally look at the drivers of variation over time in flows between cities and find evidence that it is also important to account for group differences when looking at the response of flows to changes in city

[14] This measure of cultural differences is suggested by Saiz (2010), who uses it as a proxy for a culture that is less likely to impose regulatory restrictions on housing supply.

[15] Both differences in ethnicity shares and industry shares are computed as vector distances between the origin and destination cities' vectors of shares.

[16] In particular, the results for the long sample in the first column show that physical distance and crossing state boundaries significantly increase migration costs, albeit physical distance matters less conditional on leaving the state or region. Migrants are also more likely to move to areas that are culturally similar as measured by the differences in the ethnic composition and the prevalence of non-mainline Christian denominations; and workers find it easier to move to areas with a similar industry composition. Differences in local access to the water increase migration costs.

[17] The only exception is that migrants were actually *more* likely to move between cities with greater differences in January temperatures, perhaps reflecting the large ongoing movement towards the warmer regions of the American South (Glaeser and Shapiro, 2001).

[18] This is in line with findings by Molloy et al. (2011) that interstate migration rates increase with education during the 1980-2010 period.

characteristics. In the quantitative model estimated in Section 6, I will account for this subgroup heterogeneity, as well as the effect of observed and unobserved changes in amenities when trying to estimate the location choice parameters.

As an example of the estimated effects of differences between cities on migration, consider again migration into Boise, ID: Its industry structure is much more similar to that of Los Angeles than the difference for an average city pair.[19] The reason is that both cities have an unusually large share of local employment in business services industries, such as finance and information technology, which represent 24% of year 2000 employment in Boise, and 29% in Los Angeles. In contrast, these industries only represent 14% of employment in the average commuting zone.[20] The predicted effect of this similarity in industry structure on migration, based on the estimates in Table 2, is that college-educated migration flows from Los Angeles to Boise should be 31% times larger than for a city pair with average industry structure differences.[21] Similarly, the shorter distance from Los Angeles is predicted to lead to 46% greater flows of college-educated workers to Boise, ID, than to Wichita, KS, for example.[22] Of course, part of the reason for greater migration links between particular city pairs that is not captured in this regression may be historical path dependence: migration flows in the past – whatever their cause – lead to greater family and social connections between cities, and more information about opportunities at the destination, and thereby make it a more likely destination in the future. For example, already in 1980, 7% of all people in Idaho (where Boise is by far the largest city), were born in non-neighboring California, making it the most important state of origin for non-native residents.[23]

Overall, these results suggest that the persistence of migration networks is partially explained by workers seeking out cities that are physically and institutionally close, and which have similar cultures and employment structures. This rationalizes the use of historical migration shares in the reduced form analysis as sufficient statistics for the potential of later migration spillovers between particular city pairs.

2.3 Case study: Superstar cities and migration spillovers

What do migration links look like among large cities in the U.S – and how does that affect house price dynamics? In this section, I show an example of how identifying the migration networks of a particular set of cities with particularly strong economic booms allows us to predict which other cities will experience population growth. I focus on a small set of "superstar" cities that were shown by Gyourko et al. (2013) to have had historically particularly inelastic housing supply and high housing demand, which result in high house price elasticities with regard to population growth, which means that "lower income households [are] crowded out by higher income households"(Gyourko et al., 2013).

[19] The 2-digit industry share vector distances are 0.07 for Boise-L.A. and 0.28 on average.

[20] I follow Eckert et al. (2019) in defining Business Services employment as consisting of NAICS-5 industries, that is, it includes NAICS sectors 51, 52, 53, 54, 55, and 56. The employment shares are computed from QCEW data.

[21] This is calculated as $exp(-1.291 * 0.07)/exp(-1.291 * 0.28) = 1.31$, where I use the coefficients on industry structure differences in Column 2 of Table 2.

[22] This is calculated as $exp(6.280*(-.682))/exp(6.836*(-.682)) = 1.46$, where $6.280 = ln(533)$ is the log distance in miles from L.A. to Boise, and $6.836 = ln(931)$ that from L.A. to Wichita.

[23] The question of state of birth by current state of residence was asked for the first time on the 1980 census. Overall, 50.9% of all Idaho residents were born out-of-state in 1980, and 14% of the latter were born in California. Census data is available from: `https://www.census.gov/data/tables/time-series/demo/geographic-mobility/place-of-birth-decennial.html`

As a result of these characteristics, superstar cities should be very likely to originate migration flows during the housing boom 2000-2007 that lead to spillovers to destination cities. Gyourko et al. (2013) identify superstar cities by noting for the respective 20 years preceding the years 1970, 1980, 1990, and 2000, whether a metro area was above-median in the sum of house price growth and housing unit growth ("high demand"), and in the top decile of the ratio of these two variables ("inelastic"). Superstars fulfill these criteria for at least two of the four periods considered. The resulting list of 21 superstar MSAs corresponds to 14 different 1990 commuting zones, which I use in my analysis.[24]

Identifying migration spillover cities. In order to identify cities that are likely to be impacted by migration spillovers from superstar cities, I proceed as follows: First, I restrict the search to large cities[25] Second, for each superstar city, I compute the average annual share of all IRS-reported migration outflows going to each other city for the years 1990-2000. Then, for each superstar city, I retain the Top 4 large migration destinations receiving the highest 1990-2000 outflows. Any of the retained top destinations that are not superstar cities themselves, I add to the list of "superstar spillover cities", of which I identify 12.[26] Note that both the superstar and the spillover cities were defined solely based on data up to the year 2000.

The housing boom in superstar and spillover cities. To understand how the superstar and spillover cities defined *ex ante* from pre-2000 data fared subsequently during the housing boom, Figure 3 plots average house price growth and population growth during 2000-2007 for all large CZs. The graph shows that the superstar characteristics observed by Gyourko et al. (2013) in pre-2000 data are highly persistent: all the superstar cities are in the upper left corner of the graph, which means they continue to have high house price growth and a high ratio of house price growth to population growth. The "spillover" cities are characterized by high house price growth and higher population growth than the superstar cities during the boom.[27]

The graph also highlights that it is misleading to speak of "the" U.S. housing boom of the early 2000s: While house prices experienced rapid growth on average, there is substantial cross-sectional variation, with more than a third of large cities experiencing nominal house price increases of less than 5% per year. Conversely, there are clear "boom" cities with average house price growth above 10% per year, but they almost exclusively consist of superstar and migration spillover cities: of 17 "boom" CZs, 6 are superstar cities, 10 are spillover cities, and only one is neither.[28]

[24]These 14 superstar CZs, identified by their largest city, and the CZ code in parentheses, are: Albany, NY (18600); Newark, NJ (19600); Boston, MA 20500); Hudson Valley, NY (19300); Los Angeles, CA (38300); NYC (19400); New Haven, CT (20901); San Francisco, CA (37800); Philadelphia, PA (19700); Santa Barbara, CA (38200); Springfield, MA (20800); Pittsfield, MA (20902); Providence, RI (20401); San Jose, CA (37500).

[25]Here, "large" cities are defined as having more than 0.85 M adults over the age of 21 in the year 2000 census, a cutoff chosen to ensure that they represent slightly more than 50% of U.S. adults. Of all continental U.S. CZs, 45 are "large". I exclude New Orleans as it represents a large outlier in negative population growth due to the impact of Hurricane Katrina.

[26]These "spillover cities", identified by their largest city, and the CZ code in parentheses, are: Baltimore, MD (11302); Tampa, FL (6700 Miami, FL (7000); Palm Beach, FL (7100); Wash., D.C. (11304 Buffalo, NY (18000); Phoenix, AZ (35001 Fresno, CA (37200); Sacramento, CA (37400); Las Vegas, NV (37901 San Diego, CA (38000); Seattle, WA (39400)

[27]With the sole exception of Buffalo, NY, among the spillover cities, and Los Angeles among the superstars, the two groups can be neatly separated, with the cluster of spillover cities all having higher population growth as a result of migration.

[28]That one city is Orlando, FL. Also note that, among the spillover cities, there are two outliers, which have much higher population growth (> 3% per year) than other spillover cities, but are in the lower half of house price growth in this group: Las Vegas, NV, and Phoenix, AZ. Tthese two cities are unusual in how *little* house price growth they saw during the housing boom compared to their peers. While this pair of cities is sometimes

Overall, being a likely migration spillover destination of workers displaced from superstar cities seems to predict high population and house price growth during the housing boom, showcasing the house price propagation mechanism through migraton networks that I propose in this paper.

3 Data & empirical approach: reduced-form spillovers

In this section, I describe how I measure migration exposure to shocks in other cities empirically. Moreover, I discuss the construction of the instrument for house price changes in a city's migration network. Then, I show how these shocks to a city's migration network are used to identify causal house price spillover effects between cities. Last, I introduce the data sources for the reduced-form estimation.

Spillover mechanism. The analyses below all provide evidence regarding the following proposed mechanism for spillovers of house prices between cities: (1) Some cities receive economic shocks that attract one group of workers more than another (e.g. a technology shock increasing skilled workers' wages in large cities). (2) In cities with constrained housing supply, this inflow of workers increases house prices and displaces other groups (e.g. non-college workers), who are more likely to move to cities with strong pre-existing migration links. (3) These migration-linked destination cities also experience higher house price growth as a result of increased in housing demand due to the population spillover. As a consequence of this mechanism, we would observe house prices move together between cities that have stronger migration links.

3.1 Measuring Migration Network House Price Changes

To measure the spillover effects of house price changes in other cities, I need a measure of the size of a focal city's migration network exposure to house price changes in other cities. This measure should reflect exposure to changes in other cities' characteristics through migration flows. Intuitively, migration exposure of city i to changes in house prices in some other city j depends on two variables: (1) the importance of city j as a destination for migrants from all other cities k, and (2) the degree to which redirected migrants from city k will increase inflows into city i if city j becomes more expensive. These two components capture the fact that, in a network setting, migration from city k to city i will not only depend on shocks to either of these cities, but also on changes occurring in other locations that indirectly compete with them for migrants.

In the data, the importance of city j to migrants from k corresponds to the outmigration share $\mu^{k\to j}$, while the size of inflows into i that would result from city k migrants making different location choices is captured by the inmigration share $\phi^{i\leftarrow k}$ for city i.

Putting these two terms together, the migration exposure of city i to city j house prices is

highlighted as "anomalies" (Glaeser, 2013) for experiencing high house price growth during the boom, this analysis shows that an unusually large migration shock combined with their above-median supply constraints (according to Saiz (2010)) can perhaps explain their relatively large house price response. In fact, other spillover cities, such as Washington, D.C., or Baltimore, MD, which experienced similar house price increases with much smaller demand growth and similar supply constraints, should be considered much more "anomalous" during the 2000-2007 boom than Las Vegas or Phoenix.

given by

$$\psi^{ij} = \frac{1}{\bar{\psi}_i} \sum_{k \in N} \underbrace{\phi^{i \leftarrow k}}_{\substack{\text{In-migration} \\ \text{share for } i}} \underbrace{\mu^{k \to j}}_{\substack{\text{Out-migration} \\ \text{share for } k}}, \tag{2}$$

where $\bar{\psi}_i = 1/\sum_{j:j \neq i} \sum_{k \in N} \phi^{i \leftarrow k} \mu^{k \to j}$ scales each city's relative migration exposures to sum to one. I use these migration exposure weights to construct a measure $\Delta \mathcal{P}_{it}^{NW}$ of average changes in house prices that city i is exposed to through its migration network, which is defined as

$$\Delta \mathcal{P}_{it}^{NW} = \underbrace{\sum_{j:j \neq i} \psi^{ij}}_{\substack{\text{Migration} \\ \text{exposure}}} \Delta \ln P_{jt}. \tag{3}$$

This measure captures the network link of house price growth $\Delta \ln P_{jt}$ in other cities to house price growth in city i through changes in location choices. The ψ^{ij} terms weight house price growth in other cities j by the degree to which they compete with city i for migrants, considering all possible migration network links via other cities k. Appendix E.1 shows that a similar functional form for migration exposure ψ^{ij} could also be directly derived as the implied spillover effect in a static location choice model.

In order to avoid concerns over endogenous changes in these migration exposure weights over time, Ihold the migration network weights fixed across years at baseline period values. In particular, I average inter-city migration flows over 1990-1995 reported in IRS migration data to construct the migration weights $\psi^{ij}_{\text{`90-`95}}$. This practice of holding the migration weights fixed at a baseline period level has a strong precedent in the literature using historical migration shift-share instruments (Altonji and Card, 1989; Boustan, 2010; Howard, 2020; Derenoncourt, 2019). See Appendix Section F.1 for further discussion on why I choose to construct weights this way.

3.2 Baseline network spillover specification

The baseline spillover effect that we are interested in is the effect of network house price changes $\Delta \mathcal{P}_{it}^{NW}$ defined above on house price growth in city i. I estimate the following reduced-form relationship:

$$\Delta \ln P_{it} = \alpha_i + \alpha_t + \tilde{\eta}^{nw} \Delta \mathcal{P}_{it}^{NW} + \beta' \Gamma_{it} + \tilde{\xi}_{it}^{P}, \tag{4}$$

Here, the network spillover coefficient $\tilde{\eta}^{nw}$ is a function of constant model parameters, and Γ_{it} represents observable control variables capturing characteristics of the focal city or its migration network. These covariates are discussed in detail in Section 3.5. The error term $\tilde{\xi}_{it}^{P}$ captures unobservable relative changes in other drivers of location choices in the focal city relative to its migration network, i.e. amenity and wage changes, and local differences in effect coefficients.

In this specification, we are mainly interested in the network spillover effect $\tilde{\eta}^{nw}$. If house price growth in other cites can propagate through migration networks, we would expect those cities to experience higher house price growth that are more exposed to cities with increasing house prices through their migration network. That is, positive house price growth spillovers should be reflected in $\tilde{\eta}^{nw} > 0$. Moreover, given that the postulated channel for these effects runs through migration flows that cause house price increases, we should be able to find effects in the same direction for migration flows when we replace the dependent variable by the net migration rate.

In the next section, I present an identification strategy that allows me to estimate causal spillover effects of house prices in this specification.

3.3 Construction of network instruments

The main identification concern in estimating Equation 4 arises from the fact that the house price residual $\tilde{\xi}^P_{it}$ in a location choice model is likely to be a function of focal city and migration network changes in amenities and wages.

Intuitively, if house price changes in different cities are fully offset by changes in wages or amenities, then shifts in housing cost might not cause migration, as the overall attractiveness of the city would not change. Similarly, any migration network spillover effects found may be due to a correlation in the unobserved amenity and productivity shocks in connected cities rather than causal effects of one city on another. In general, there may be omitted variable bias if $Cov(\Delta\mathcal{P}^{NW}_{i,t-1}, \tilde{\xi}^P_{it}) \neq 0$. To address this concern, I develop a network instrumental variable approach for network price changes, which combines city-level Bartik (1991) shift-share instruments common in the house price literature with the network structure of migration flows.

City-level "Bartik" instruments. In standard location choice models, housing demand is a positive function of local labor productivity shocks – and the size of the house price response to these shocks will depend on the local house price elasticity (Glaeser et al., 2008). This insight motivates the use of shift-share shocks exploiting exogenous local wage changes, in combination with housing supply constraints, as instruments for house prices. I construct plausibly exogenous local wage shocks B_{it} by combining the local wage bill share $\tilde{\omega}_{\iota,i,t_0}$ of workers in 3-digit NAICS industries ι in a baseline period t_0 with national wage growth $\Delta \ln W^{US}_{\iota,-i,t}$ in that industry (as a proxy for industry wage trends) in the form

$$B_{it} = \sum_{\iota} \tilde{\omega}_{\iota,i,t_0} \Delta \ln W^{US}_{\iota,-i,t}.$$

In order to minimize bias from endogeneity in the local industry exposure – which might result from auto-correlated national industry shocks (Goldsmith-Pinkham et al., 2018; Jaeger et al., 2018) – I fix industry exposure shares at their 1990 level.[29] Moreover, the industry averages of log wage growth are computed as leave-one-out measures to avoid mechanical correlation between the national trend estimate and city i wages (Borusyak et al., 2020). See Appendix Section F.2 for more details on the construction of different shift-share shocks used in the reduced-form analysis and the structural estimation.

As an example of what might be driving common national industry trends, Eckert et al. (2019) note that the national wage growth in "skilled scalable services" like finance and communication industries in the last decades has been an important contributor to wage growth in U.S. cities where these industries are prevalent.[30] This trend, in part caused by a decline in communication costs, is an example of the identifying variation underlying the industry shift-share instruments .

To capture plausibly exogenous heterogeneity in the effect of this local exposure to national wage shocks on house prices, I also interact B_{it} with a measure of local land unavailability for construction x^{land}_i from Lutz and Sand (2019).[31] This measure captures geographic constraints

[29] Note that in the structural estimation I update the baseline period and use year 2000 industry shares instead, as the quantitative estimation only includes data for 2005 - 2017.

[30] See Appendix Figure A12 for a plot of wage trends by industry sector over time.

[31] These are comparable to the Saiz (2010) land availability measures commonly used in the literature. Lutz and

to marginal housing construction, which would be expected to increase the slope of the housing supply curve, and thereby increase the responsiveness of house prices to the wage shocks.[32]

Network instruments. Assuming that the shift-share "Bartik" shocks just defined represent plausibly exogenous shocks to local house prices in each city, a weighted average of these shocks occurring in *other* cities in a focal city's migration network should capture its exposure to exogenous house price changes in other cities. Put differently, those cities that have strong migration connections to locations that are more exposed to Bartik shocks will see greater exogenous house price growth in their migration network – which implies a greater spillover effect on their own house prices. Variation across cities in their network exposure to other cities' Bartik shocks can therefore identify the house price spillover effect $\tilde{\eta}^{nw}$.

This intuition is illustrated in Figure 4: Consider a simple migration network between Los Angeles, Boise, Boston, and Portland, where there are large migration exposures (thick solid arrows) to Los Angeles for Boise, and to Boston for Portland, while migration links across these pairs are negligibly small (dotted arrows). Now, if a wage shock B_{LA} interacts with the relatively constrained housing supply x^{land}_{LA} in Los Angeles, driving up house prices, this would be expected to cause greater migration flows from Los Angeles to Boise, but only negligible flows to Portland. As a result, house prices in Boise increase, all else equal, while Portland is mostly unaffected. In the context of this example, the network IV approach uses the degree to which Los Angeles shocks lead to greater house price changes in Boise than in Portland to identify the size of causal spillover effects.

The instruments for network house price changes are constructed analogous to $\Delta\mathcal{P}^{NW}_{it}$ as migration exposure-weighted averages of network Bartik shocks and their interactions with land constraints:

$$NWP^{B}_{it} = \sum_{j:j\neq i} \psi^{ij}_{\text{`90-`95}} B_{jt}. \tag{5}$$

$$NWP^{Bx}_{it} = \sum_{j:j\neq i} \psi^{ij}_{\text{`90-`95}} B_{jt} \cdot x^{land}_{j}, \tag{6}$$

where the migration network weights $\psi^{ij}_{\text{`90-`95}}$ are as defined in Equation 2. Intuitively, the instruments NWP^{B}_{it} and NWP^{Bx}_{it} capture the predicted exogenous component of migration-weighted house price growth in city i's migration network.[33]

Sand (2019) build on his methodology to expand the number of covered cities and, among other things, improve the measurement of land availability for overlapping city areas and coastal locations.

[32] An influential critique by Davidoff et al. (2016) of research using housing supply constraints like land availability as instruments for house prices has noted that supply constraints may be invalid instruments due to their correlation with housing demand shocks. That critique is not applicable here: Most importantly, my main analyses identify effects off variation in Bartik shift-share shocks over time, and the supply constraints only matter in so far that they form part of the exposure term in the shift-share instrument construction. As Borusyak et al. (2020) show, it is not necessary for these exposure terms (and therefore the supply constraints) to be exogenous with regard to local unobservable shocks, as long as the shifter in the form of a national industry trend is exogenous. The precise sense in which this exogeneity condition needs to hold here is detailed in Section 3.4. In addition, my main analyses control for commuting zone fixed effects that eliminate any constant effect of supply constraints on house price growth. Beyond that, I show that my analysis is robust to controlling for a number of different local characteristics that should at least in part control for any local dynamics caused by supply constraints. Most importantly, I control directly for the interaction of the focal city's shift-share shock with local supply constraints (see Section 3.4) – and the identifying variation is only coming from the residual effect of *other* cities' supply constraints interacting with shift-share shocks. For these reasons, I believe the Davidoff et al. (2016) critique of supply constraints used in IV settings is not relevant in my setting.

[33] Analogous to the construction of the wage shocks B_{it}, and consistent with the definition of $\Delta\mathcal{P}^{NW}_{it}$, I again

There is by now an extensive literature that instruments for local house prices using shift-share measures of local exposure to exogenous national productivity trends (Saks, 2008; Saiz, 2010; Paciorek, 2013; Mian et al., 2013; Guerrieri et al., 2013; Beaudry et al., 2014; Diamond, 2016). The innovation in this paper consists of applying those city-level shocks to *other* cities to identify exogenous house price changes in the focal city's *network* and their effect on the focal city's house prices through a spillover channel, rather than looking at the direct effect of the Bartik wage shocks on the focal city. As far as I know, there are few comparable papers that use a network approach to shift-share identification. The most similar ones include Bartelme (2018), who uses Bartik shocks to nearby cities as an instrument for a city's market access in a trade gravity model. Similarly, Baum-Snow et al. (2019) use shift-share labor demand shocks to other neighborhoods that are accessible by commuting to identify the effect of changes in "resident market access" on life outcomes. However, I am not aware of any other work using migration links and shift-share shocks to other cities to identify spillover effects.

3.4 Network instrument identification

Industry-level reformulation of network IV. In order to make the discussion of the identification in my network IV approach more transparent, I first extend the results in Borusyak et al. (2020) to show how the distribution of industry wage shocks determines the spillover estimate. As I show in Appendix Section E.2, the IV estimator of the spillover coefficient can be written as

$$\hat{\tilde{\eta}}^{nw} = \frac{\sum_{t=1}^{T} \sum_{\iota=1}^{N_{ind}} s_\iota g_{\iota t} (\overline{\Delta \ln P_{it}})_{\iota t}^{\perp}}{\sum_{t=1}^{T} \sum_{\iota=1}^{N_{ind}} s_\iota g_{\iota t} (\overline{\Delta \mathcal{P}_{it}^{NW}})_{\iota t}^{\perp}},$$

Here, I have combined the industry exposure and migration network structure into a weight $s_\iota = \frac{1}{N} \sum_{i=1}^{N} \sum_{j:j \neq i} \psi^{ij}_{\text{‘90-‘95}} \tilde{\omega}_{\iota,j,\text{‘90}}$ that summarizes the average across cities i of their migration network exposure to industry ι. The notation $\bar{\nu}_t$ denotes an average across cities, weighting them by their relative exposure to industry ι, and $g_{\iota t}$ is log national wage growth in industry ι. This estimator is equivalent to running an industry-level regression of weighted average house price growth (with weights given by industry exposure through migration networks) on weighted network house price growth, instrumenting for the latter with national industry wage growth.

This rewriting of the estimator in this way at the industry level clarifies the identifying variation underlying the network IV estimate: We can think of the spillover effect estimation as identifying the spillover effect from the degree to which the covariance in industry wage growth shocks with house price growth in the cities most exposed to the industry (the numerator) is higher in the cities that are "treated" in the form of having house price changes in their migration network that vary with industry wage shocks (the denominator). For instance, if cities with migration connections to technology hubs (e.g. Boise, ID) see higher house price growth precisely when tech hub house prices rise (e.g. in San Francisco) as a result of national trends in knowledge industry wages, then this variation allows us to infer a causal positive spillover effect.

This rewriting of the network IV estimator in the form of industry-level shocks then allows me to formulate the exclusion restriction of the network approach as follows:

Network IV exclusion restriction: *If the network instrument NWP_{it}^{B} (or NWP_{it}^{Bx}) is relevant and mild regularity conditions hold (the variance matrix of control variables has full rank, and the*

hold the migration exposure shares fixed at a baseline period level, here consisting of average 1990-1995 migration shares.

covariance matrices of instruments and residuals with controls are bounded and exist), then the IV estimate of the spillover effect $\hat{\hat{\eta}}^{nw}$ is consistent if and only if

$$\sum_{t=1}^{T}\sum_{\iota=1}^{N_{ind}} s_\iota g_{\iota t}\bar{\xi}^{P,\perp}_{\iota t} \to^p 0.$$

Here, $\bar{\xi}^{P,\perp}_{\iota,t}$ is the error in the house price growth regression in Equation 4, residualized with regard to the control variables Γ_{it}, and averaged over cities, weighting them by their migration network exposure to industry shocks.[34]

This condition shows that, for the network IV estimate to be consistent, industry wage shocks cannot be systematically higher for those industries that have a systematically larger migration network impact on cities that are experiencing large unobserved house price shocks, conditional on control variables.

As Borusyak et al. (2020) show, this identification allows for a city's migration network to be endogenously determined – it only requires the national industry trends over time to be exogenous in the sense defined in the exclusion restriction. This would be invalidated, if, for example, cities that experience more migration flows from cities that specialize in the booming tech industry are also systematically experiencing greater idiosyncratic house price movements in a way that is not captured by their own exposure to knowledge industries or any other included control variables.

"Double Bartik" control variable. As I am using the industry structure of other cities in the same migration network to instrument for network house price shocks, there may be a concern that the network house price instruments are correlated with focal city i industry shocks if industry structure is correlated across cities that share migration links. Note that this concern is supported by the significant coefficient on industry structure in the migration cost determinants analyzed in Section 2.2: Table 2 showed that migration costs appear to be lower among cities with similar industry structures, making them more likely to have strong migration links.

To address this issue, I include focal city i's direct Bartik shocks B_{it} and $B_{it}x^{land}_{it}$ as control variables in the regression, an approach which Chodorow-Reich and Wieland (2020) titled the "double Bartik" method. Controlling directly for the focal city's own Bartik shocks effectively orthogonalizes the network instrument with regard to any direct city i exposure to national industry wage shocks. To make this concrete, consider again the Boise, ID, example: While Los Angeles and Boise have a relatively similar industry structure, it is not identical. For example, L.A. is more exposed than Boise to industries in the "information" industry sector that includes motion pictures and broadcasting (NAICS code 51), which saw real wage growth nationally in 2000-2017 that was 15% higher than average.[35] This difference in exposure to shocks means that Los Angeles will see some increases in local house prices that are not correlated with Boise's direct exposure to industry wage shocks – and these effects will be larger due to L.A. being in the top quartile of land-constrained cities. These exogenous house price increases in L.A. that are driven by industry shocks *not* affecting Boise are then used to identify the spillover effect on house prices in Boise due to its exposure to migration outflows from L.A.

Note that it is not necessarily an issue that the Bartik industry shocks might be an imperfect proxy for unobserved (and potentially correlated) actual industry shocks that are affecting both a

[34]See Appendix Section E.2 for details on the derivation of this expression.

[35]Real wage growth (corrected for CPI-U growth) in NAICS 51 was 27% over the 2000-2017 period, while the employment-weighted average of real wage growth across sectors was 12%. See Appendix Figure A12 for an illustration of these industry wage trends.

focal city and its migration network. As the analysis above showed, if there is measurement error in the industry exposure terms that affects local house prices, the exclusion restriction is only violated if these measurement errors and unaccounted-for industry similarities are systematically correlated with migration network exposure to growing industries, *and* are not captured by the included control variables, such as local Bartik wage shocks.

3.5 Additional covariates and robustness checks

In addition to the IV strategy described above, I also try to address any potential omitted variables bias by explicitly controlling for covariates – informed by a location choice model – that might confound the house price spillover effect. Moreover, I explore different specifications to ensure the robustness of the results. These variations are detailed below.

National house price trends. An extensive literature suggest that house prices might co-move in different cities due to national house price trends, for instance caused by changes in mortgage rates or business cycle dynamics as a result of credit supply changes (López-Salido et al., 2017; Mian et al., 2019). I control for this possibility by including year fixed effects in the regressions.

Regional trends. There may also be regional differences in house price trends, for instance due to the timing of settlements in the historical evolution of the U.S., resulting in higher population mobility in general due to differences in the "rootedness" of local populations, as Coate and Mangum (2019) argue. Moreover, underlying house price trends may be caused by long-run shifts in preferences for regional amenities, such as the increased attractiveness of the sunbelt as a result of air conditioning (Glaeser et al., 2001). To avoid confounding my analysis of house price-driven migration spillovers with these trends, I allow for U.S. Census region-specific time trends and city fixed effects in house price growth.

Migration access. The static location choice model in Appendix E.1 suggests that the size of migration flows to and from a city should depend on the population size of cities in its migration network. That is, a given change in attractiveness of a city should result in a different migration flow if "migration access" in the form of the potential migrant population changes – analogous to the role played by market access in trade models (Donaldson and Hornbeck, 2016; Allen and Donaldson, 2018). To control for this size effect, I construct the empirical counterpart to migration access $\mathcal{M}_{i,t-1}$ implied by the model as $\mathcal{M}_{i,t-1} = \sum_k \phi^{i \leftarrow k}_{\text{`90-`95}} \Delta \ln L_{k,t-1}$, and include it in the regressions.

Wage effects. The attractiveness of a city to migrants is not just a function of housing costs, but also of income. In fact, if a shock to labor demand creates opposite and offsetting wage and housing cost effects, we might see a city becoming more expensive but without any change in migration as real wages remain constant (Moretti, 2013). Therefore, it may be important to control for wage changes in the migration network. To do so, I construct the wage change analogue of the network house price growth term as

$$\Delta W^{NW}_{it} = \sum_{j: j \neq i} \psi^{ij}_{\text{`90-`95}} \Delta \ln W_{jt},$$

where W_{jt} is computed as the salary income per capita in the CZ reported to the IRS. I include ΔW^{NW}_{it} as a control variables in my baseline specification.

Neighboring CZ spillovers. There is a possibility that local shocks could affect neighboring commuting zones directly through changes in commuting patterns or overlapping housing markets – without a need for a migration channel. For instance, shocks to the New York financial sector may affect house prices in some parts of Northern New Jersey directly through the incomes of long-distance commuters. I control for this possibility by excluding from any network measures commuting zones that are too close to one another in terms of physical distance. In the most stringent baseline specification, I will exclude any CZs that contain counties of which the county center is less than 150 miles away from *any* county center in commuting zone i. This ensures that the estimated house price spillovers operate at a distance that requires long-distance migration, such that a worker could not retain their local job when moving. Note that this inverts the mechanism considered, for example, in DeFusco et al. (2018), who focus on spillovers between nearest-neighbour cities.

Endogenous lagged house price growth. To explore the possibility that extrapolation plays a role in house price formation, I also consider specifications that include lagged house price growth as an independent variable. However, when I include lagged house prices $\Delta \ln P_{i,t-1}$ in city i itself in the network spillover estimation, they are likely to be correlated with the contemporaneous house price error terms. Thus, the coefficient on these variables would be biased and would in turn lead to bias in our network spillover estimate. To prevent this issue, I instrument for any lags in house price growth using the interaction of the Bartik industry shock, for the corresponding lead or lag period, and local land availability. Moreover, as the network spillover effect would also operate in any past period, I also include $\mathcal{P}^{NW}_{i,t-1}$ in the extrapolation specification and instrument for this variable using two lags of the network instruments.

3.6 Data

As there is currently no long panel of annual U.S. migration data available that has information on different worker types, as well as sufficient sample size to accurately capture the network of city-level migration, I will limit the reduced-form analysis to the CZ-by-year level – implicitly treating workers as homogeneous. In the full quantitative estimation in Section 6, I will relax this assumption and allow for differences between college- and non-college-educated workers. This section documents the main data sources to implement the reduced-form network IV approach described above. All data is crosswalked to 1990 commuting zones, aggregating county-level data using a crosswalk provided by Autor and Dorn (2013).[36]

Migration flows. I compute migration flows between commuting zones from county-to-county migration flows provided by the Internal Revenue Service (IRS) for 1990-2017. The IRS uses changes in the zip codes on individual income tax returns that were filed for the previous tax years to infer moves. Based on this methodology, the IRS is able to assemble a data set of the annual movement of tax returns and the number of exemptions – which correspond closely to the number of people – across counties, covering $> 90\%$ of the U.S. population.[37] This mobility data should correspond closely to the actual movement of people across counties - with small caveats (Gross, 2003).[38] I use this data to compute CZ-to-CZ flows and the share of people not moving in

[36]Available online at https://www.ddorn.net/data.htm.

[37]For instance, for the fiscal year 2014, the IRS received 149 million income tax returns - which correspond to tax households - on which 291 million personal exemptions were claimed (Source: https://www.irs.gov/uac/soi-tax-stats-individual-income-tax-returns), which compare to ca. 319 million people in the U.S. in 2014.

[38]The small number of people who do not have to file income tax returns is not captured in this data, and in that number low-income people and the elderly tend to be overrepresented. At the same time, a small number of tax returns that are filed late - after September of the filing year - are not captured. The latter returns are usually granted extensions due to their complexity, which correlates with high income, so the mobility data may under-represent very high income returns as well. Moreover, changes in marital status may lead to a failure to

each year, omitting flows in and out of New Orleans from any totals.[39] The baseline in-migration weights $\phi^{i \leftarrow k}_{\text{'}90\text{-'}95}$ are calculated as the average inflow from city k to i in each year 1990-1995, as a share of the average population in city i in those years. Similarly, the share of people staying in a city $\mu^{ii}_{\text{'}90\text{-'}95}$ is the average number of people (tax exemptions) not moving away from city i as a share of the previous year's population – also averaged over 1990-1995.

For the construction of the city-to-city migration network, I only include continental U.S. commuting zones, i.e. not including flows to and from Hawaii, Alaska, Puerto Rico, or other U.S. territories. For migration totals, for instance to calculate gross inflows, outflows and net migration from each city, I take total migration to *any* destination and subtract foreign migration to obtain a measure of total domestic migration flows, albeit not restricted to the continental U.S., which cannot be separately identified in total flow data.

House price growth. An annual panel of house prices is obtained from the Federal Housing Finance Authority (FHFA) at the county level for 1990-2017. Different counties have data starting in different years, but by 1990 house price indices are available for most of the counties in the U.S. The data are repeat-sales indices, so while the levels are not comparable cross-sectionally, they can be used to construct the house price growth terms relevant for the reduced-form analysis.

I first compute county-level log changes in house prices and then aggregate these growth rates to commuting zones as weighted averages, with weights derived from county populations in 2000. Moreover, to avoid bias from small counties entering the sample in later years as FHFA coverage improves, I fill in the county-level growth rates with state-level house price growth for CZ-years where less than 10% of the CZ population are in counties with available house price data. This imputation is a minor issue as CZs with more than 50% of their population being covered by FHFA reporting constitute 95% of the continental U.S. population by 1990 and 99% by the year 2000.[40]

The resulting balanced panel of 1990-2017 house price growth for all CZs is the dependent variable used in the reduced form regressions.

Housing construction permits. New housing construction permits by commuting zone are constructed from county-level counts of permits available from the Census Building Permits Survey for 1990-2018, which I aggregate to 1990 CZs. Some counties do not report permit numbers. Thus, in cases where not all component counties of a CZ are reporting, I scale the reported numbers in proportion to the reporting counties' share of the CZ population to make up for the missing share, using county populations in 2000 as weights. To minimize the impact of any scaling, only CZ-years with counties reporting that correspond to at least 80% of the CZ population are included in the analysis.[41]

match returns, as for joint filers only the primary taxpayer's social security number is captured, so that after a divorce, for instance, one of the former couple's social security numbers will not be associated with a corresponding county of residence for tax purposes for the prior year and thus gets omitted from the data set. It is also important to note that the IRS does not name counties if the pairwise flow consists of less than 10 tax returns, so small migration flows are censored. However, these flows are contained in the totals.

[39]The reason for excluding flows to and from New Orleans here and in all other analyses is that the city's loss of the majority of its population as a result of Hurricane Katrina in 2005 represents a large outlier that would distort all migration-related analyses and is mostly unrelated to the deliberate equilibrium location choices that I focus on in this paper.

[40]In unreported regressions, I have also verified that the main results of the reduced-form analysis are robust to omitting all CZs that did not have more than 50% FHFA reporting coverage for their population for the entire 1990-2017 sample period.

[41]This is a minor issue, as <8.6% of county-years are part of CZs that have <99% of their populations accounted for. In any year, at least 717 CZs have at least 80% of their total population fully covered by permit reporting.

Wages. I use IRS tax data on total salary income by county to compute average salary income per capita for each commuting zone as a measure of average wages.

Mortgage lending volume and loans. In order to measure the effect of migration on mortgage market activity, I use loan-level data on mortgages provided under the Home Mortgage Disclosure Act (HMDA). For each year and borrower county, I retain all the loans that are originated for home purchases (excl. refinancing and other loan purposes) and aggregate total lending volume in dollars, as well as the count of new originated purchase loans to the CZ level.

Vacancy data. I obtain vacancy rates from the Census Housing Vacancy Survey and compute changes in average vacancy rates for homeowners and renters between a baseline period (1991-1997) and the 2000-2007 and 2012-2017 boom periods in my analysis. This comparison can be computed for 98 CZs, which have data for all time periods. I group these cities into those with above- and below-median housing supply constraints based on their share of land unavailable for construction (Saiz, 2010; Lutz and Sand, 2019).

4 Reduced-form estimates of house price spillovers

In this section, I first present the baseline results on house price spillovers between cities. Second, I show that the expected effect pattern holds for net migration to be an important transmission channel, and that other city-level variables behave in line with housing demand increasing as a result of migration. Third, I demonstrate that migration spillovers are quantitatively important by estimating their ability to explain the cross-sectional variation in house prices in the run-up to the housing boom of the 2000s, and by showing that migration links are good predictors of house price correlations between city pairs.

4.1 Migration spillover effects on house prices

4.1.1 Baseline effect estimates

Recall that I am estimating equations of the form

$$\Delta \ln P_{it} = \alpha_i + \alpha_t + \tilde{\eta}^{nw} \Delta \mathcal{P}_{it}^{NW} + \beta' \Gamma_{it} + \xi_{it}^{P},$$

where Γ_{it} represents various control variables capturing characteristics of the focal city or its migration network. Moreover, in the baseline specifications I will only allow for CZs to be included in one another's migration network, if they are at least 150 miles apart.

OLS baseline estimates. The effect of estimating Equation 4 using OLS without any of the control variables is shown in column 1 of panel A in Table 4. I find a significant positive coefficient of house price changes in a city's migration network on the focal city's house price growth. The other columns of the table sequentially add in the control variables discussed in Section 3.5: In column 2, I control for city fixed effects, as well as national trends. In column 3, I additionally control for regional time trends. In column 4, I add the lagged migration access. In column 5, I control for average wage changes in the migration network. In column 6, I add the "double Bartik" controls for focal city labor productivity shocks, as well as their interaction with land availability. The coefficient on the network house price changes is remarkably stable: even after adding in all the control variables in column 6, the coefficient size changes by less than 10% relative to the estimate in column 1.

To interpret the effect sizes, note that the network house price term scales house price changes in other cities by migration weights $\psi^{ij}_{\text{'90-'95}}$, which sum to one. As a result, the coefficient can be interpreted as the pass-through of average house price changes in city i's migration network. That is, an average 10 ppt house price change in all other cities in the migration network correlates on average with a contemporaneous change in house prices growth in city i of about 1.4 ppt (based on the OLS estimate in Column 6), after accounting for the included control variables.

Network IV first stage. Before proceeding to the causal effect estimation, I verify that the network labor productivity shocks affect network house price changes in the expected manner – the "first stage" of the estimation. That is, we are interested in the effect of the instruments NWP^{Bx}_{it} and NWP^{B}_{it} on the migration-weighted index of actual house price changes in each city's network across all years 1991-2017, after residualizing all variables with regard to the full set of control variables. The corresponding first-stage coefficient estimates are shown in Column 2 of Table 3. Note that the relationship is upward sloping for the land constraint interaction with the wage shock, and downward sloping for the wage shock alone.[42] This means that the more constrained the migration network of a city is, the more positive the effect of shocks on network house prices is. The same relationship is also shown visually in Appendix Figure A3 as binned scatter plots of the instruments NWP^{Bx}_{it} (left panel) and NWP^{B}_{it} (right panel) with the residualized migration-weighted index of actual house price changes in each city's network.

The first-stage F-statistics for the instruments, corresponding to the predictive relationship shown in Figure A3, are reported at the bottom of the main IV results in Table 4, and they are large in all specifications based on conventional thresholds.

Network IV baseline estimates. Panel B of Table 4 shows the IV-GMM estimates using the network productivity shocks to instrument for network house price changes. The IV coefficients are positive, statistically significant and sizeable in all specifications. Column 6 of Table 4 shows that the estimate of the network house price effect when we include the full set of control variables is about 60% larger than the corresponding OLS estimate, suggesting that there was some omitted variable bias. This negative bias in the OLS could for instance be explained by the house price increases being in part caused by positive shocks to city's amenities which will reduce the ouflows of migrants to other cities in response to the higher housing costs, inducing a more negative correlation in house prices which the IV corrects.

The estimated causal spillover effect size implies that a 10 ppt average increase in house price growth throughout a city's migration network has a causal effect of about 2.3 ppt higher house price growth in the city itself. This means that a one standard deviation increase in exposure to house price growth in other cities (see summary statistics in Appendix Table A1) leads to a 0.2 standard deviation change in city i house prices. Note that these reduced-form effects do not necessarily capture the full general equilibrium effect of shocks, as a shock that spills over to city i could then have secondary spillover effects to cities that are exposed to migration as a result of city i house prices.

[42] The intuition for this downward slope is as follows: the raw relationship is positive - see Column 1 of Table 3. However, industry wage shift-share shocks increase both wages and house prices, and I am interested in the out-migration driven by increases in housing cost net of any increase in wages. Once I include the appropriate wage controls, I find the effect predicted by theory: controlling for wage changes, migration networks with lower supply constraints must have lower house price growth for a larger labor demand shock. That is, when we observe cities with a large and a small positive shock, but with the same wage effect, this suggests that house price growth must have been lower in the large shock city, such that real wages went up more. However, it is always the case that when a larger shock coincides with higher supply constraints, then house prices increase more even holding wages constant.

Dynamic effect methodology. To understand the time patterns of these spillover effects, I study the effect of period t shocks on city i in future periods. I estimate IV forecasting regressions that correspond to local projections with external instruments (Jordà, 2005; Stock and Watson, 2018) of the form

$$Y_{i,t-1+h} = \alpha_i + \alpha_t + \tilde{\eta}_h^{nw} \Delta \mathcal{P}_{it}^{NW} + \beta' \Gamma_{it} + \xi_{i,t-1+h}^P, \tag{7}$$

where the dependent variable $Y_{i,t-1+h}$ can be a flow variable, such as period $t-1+h$ house price growth, or a cumulative variable, such as the total effect on the level of house prices $\sum_{k=1}^{h} \Delta \ln P_{i,t-1+h}$ of period t shocks, or the effect on the level of population. The vector Γ_{it} includes the same additional period t control variables as the static model (see Section 3.5), and I am again instrumenting for the time t network price growth shock using network shocks. The coefficients $\tilde{\eta}_h^{nw}$ now represent the impulse response in period $t-1+h$ of the shock. That is, the contemporaneous impacts correspond to $h=1$, the impact on the dependent variable in the year after the shock is $h=2$, and so on.

When reporting dynamic effects, I will cumulate impulse responses on house price growth and net migration to obtain total house price level and population level effects over time. However, effects on per-period flow variables like permits or mortgage lending will be reported as per-period effects, as changes in these flow variables do not cumulate.

Dynamic house price effects. The estimated time pattern of IV estimates of period t shock effects on the level of house prices is shown in Panel (a) of Figure 5. The graph shows the cumulative effect of period t network house price changes on focal city house prices at different horizons h, where the $h=1$ estimates correspond to contemporaneous price effects shown previously in Column 6 of Table 4. As the graph shows, the spillover effects increases for two periods past the initial impact, and then levels out. The long-run effect of a spillover shock in the fifth year (four years after the initial shock) is 0.63 and statistically significant.[43] This means that an exogenous house price change of average size 10 log points in city i's migration network causes a 6.3 log point change in city i house prices in the long run, about 2.7 times the initial impact in period t. This highlights that effects operating through a migration channel may have a delayed impact as frictions prevent workers from moving instantaneously. However, the shape of the dynamic effects graph suggests that most of the long-run spillover effect has been realized 1-2 years after the initial shock.

4.1.2 Additional house price effect specifications and robustness checks

This section shows that the baseline result of significant causal spillover effects on house prices is robust to accounting for heterogeneity in land constraints for the affected cities, as well as allowing for autocorrelation in house prices. Moreover, I show that including the full network of migration links in the analysis – rather than excluding commuting distance neighbors – if anything, strengthens the result.

Heterogeneity in spillover effects Above, I established that house price shocks in a city's migration network increase that city's house prices. However, the effect of these spillover demand shocks on each city might differ based on the degree to which it has elastic housing supply (Glaeser et al., 2008). To test whether there is heterogeneity in this regard, I allow the network spillover effect to vary by city in the form

$$\eta_i^{nw} = \psi_L + \psi_L^{land} x_i^{land},$$

[43]See Column 1 of Appendix Table A3, which shows the numerical estimates for the long-run effects plotted in Figure 5.

where x_i^{land} is the city share of land unavailable for construction, which proxies for geographic constraints to housing supply (Saiz, 2010; Lutz and Sand, 2019). The coefficient on the interaction term should be positive if the housing supply function in response to spillover shocks is steeper for cities that are more land-constrained.

The results are shown in Column 1 of Table 5. Here, I include the same control variables as in the baseline specification, with the exception of the focal city wage shock inter action with land constraints: Including this interaction as a control variable would remove most of the heterogeneity in responsiveness to shocks that we are interested in. That is, I cannot precisely identify the differential effect of network spillovers, while holding the differential impact of direct wage shocks constant – and it would be difficult to interpret such an estimate.[44]

The coefficient on the constraint interaction term is positive and significant at a 1% level. That is, cities that are more constrained experience stronger house price responses to network spillovers. To interpret the magnitudes of this heterogeneity, note that it implies that a one standard deviation change in the unavailable land share of ~ 22 ppt would increase the short-run spillover elasticity by about $0.22 * 0.42 \approx 0.09$, or about 39% of the original contemporaneous spillover estimate of 0.23 (see Column 6 of Table 4).

The house price effects over time when allowing for heterogeneity are shown in panels (c) and (d) of Figure 5. The results show that constraints increase the house price response even more in the long-run than in the short-run. This is in line with the intuition that, while both constrained and unconstrained cities experience congested construction and housing markets in the short run, land constraints additionally limit cities' ability to adapt in the long run by expanding their housing stock.

Allowing for house price extrapolation. It is possible that the identified house price effects in some period t do not only reflect shocks in period t (or in previous periods) but also are autocorrelated for other reasons. For example, if homebuyers are myopic and infer underlying trends in housing demand from past changes (see, e.g. Glaeser et al. (2017)), then house price changes may show extrapolative dynamics where house price changes in one period cause further house price changes in future periods. This raises a potential issue for identifying the size of the spillover effects: Part of the period t house price changes may be the result of past periods' house price changes, which would mean that the size of our contemporaneous effect estimates might be biased.

To explore how such autocorrelation might affect the results, I amend the dynamic spillover specification to be

$$\sum_{k=1}^{h} \Delta \ln P_{i,t-1+h} = \alpha_i + \alpha_t + \tilde{\eta}_h^{nw} \Delta \mathcal{P}_{it}^{NW} + \beta_h' \Gamma_{it} + \tilde{\eta}_h^{x} \Delta \ln P_{i,t-1} + \tilde{\eta}_h^{nwx} \Delta \mathcal{P}_{i,t-1}^{NW} + \tilde{\xi}_{i,t-1+h}^{P,cum},$$

where I have added terms that allow house price changes in response to a period t shock to also depend on lagged house price changes in city i given by $\Delta \ln P_{i,t-1}$ and lagged network house price changes $\Delta \mathcal{P}_{i,t-1}^{NW}$. These lagged terms are likely to suffer from the same endogeneity issues as contemporaneous network house price growth. Therefore, I instrument for them in an analogous way, adding two lags of the network shock instruments (with and without land constraint interactions) and also the focal city wage shock interacted with its land constraint, to the set of instruments.

The estimated contemporaneous effect $\tilde{\eta}_1^{nw}$ under this specification is shown in Column 2 of

[44] In fact, when I control for this interaction, only the uninteracted effect coefficient of network shocks remains significant, while the interaction is estimated not to have an additional effect, or a weakly negative effect.

Table 5, and the full curve of dynamic effects is shown in Panel (b) of Figure 5. The estimated contemporaneous spillover effect is not significant, but the cumulative effect becomes significant by the second year after the shock. That is, some of the effect in the initial year from the long-spillover shock is attributed to extrapolation from past shocks. However, the long-run effect is very similar to that found without allowing for autocorrelation, suggesting a long-run passthrough from the migration network of 0.63 (see Column 3 of Appendix Table A3) that is almost identical to the baseline long-run effect.

The role of migration distance. One concern with the baseline specification may be that I am restricting the migration network for each city to not include other cities that are less than 150 miles away. As the size of migration flows is in part driven by distance (see the estimates in Table 2), this may preclude a substantial share of migration flows from being considered in these estimates. Moreover, as distance has a differential effect on different education groups,[45] the migration network constructed from longer-distance flows may select for the effect of particular workers' migration.

To analyze the role of expanding the migration networks to include short-distance flows, I repeat the baseline effect estimation of Column 6 in Table 4 for migration networks defined at different distances. The results of this exercise are shown in Appendix Table A2. The minimum distance of included cities declines from the right to the left in the table, with Column 1 showing estimates for the full migration network. The spillover effects for short-distance migration networks become larger as more cities are included, with the full network IV coefficient in Column 1 being three times as large as the baseline estimate in Column 4. Moreover, the sample size increases as the migration network is expanded because some smaller cities without long-distance migration flows are represented in the sample. Notably, the strength of the instrument reflected in the 1st-Stage F-statistic declines in spite of the larger sample size when we include all cities in the network.

These effects are aligned with the original rationale for excluding short-distance migration: spillover effects between neighboring cities may operate through a variety of channels that do not require migration, as their labor markets and housing markets may partially overlap. Therefore, the effect of interest which operates through migration spillovers is more precisely identified at a distance that ensures that workers would have to move residences in order to respond to economic shocks in other cities. While the short-distance effects would imply a larger role for spillovers, I will therefore conservatively rely on the long-distance migration network estimates in Column 4.

4.2 Testing the Spillover Mechanism

If migration flows and their impact on housing demand are the mechanism through which house price changes propagate, we should be able to observe spillover effects on other characteristics of city i in line with the house price effects documented in the previous section. This section documents that house price spillovers are accompanied by large effects on migration and credit variables in affected cities that support the proposed spillover mechanism.

4.2.1 Spillover effects on population

For the spillover effects to operate by increasing net migration flows as a result of other cities becoming more expensive, the population in the affected city should increase when house prices increase as a result of the shocks. I can test this prediction of the mechanism by estimating the

[45] Compare Columns 2 and 3 in Table 2.

cumulative network spillover effect on population as a result of net migration in the same way as we estimated the cumulative house price effects, but with city i's cumulative net migration rate in the IRS data as the dependent variable.[46]

The estimates of the contemporaneous effect of network house prices on population changes as a result of net migration are shown in Column 3 of Table 5. In line with the model prediction that the network house price effects are transmitted through migration, I find that the network effect on population is positive and significant. The coefficient estimate for the contemporaneous network house price effect on net migration is 0.05 for the unrestricted migration network. This suggests that a 10 log point increase in house price growth in a city's migration network would drive a contemporaneous increase in the net migration rate of 0.5 ppt.

Long-run effects on population The dynamic effects on population (cumulative sums of log population changes from net domestic migration) are shown in Panel (a) of Figure 6. The graph shows that a 10 log point increase in network house prices in period t leads to an increase in local population by 1.9 log points 4 years after the shock.[47] This means that the long-run population effect almost quadruples the contemporaneous effect shown in Table 5. The dynamic effects show that most of this increase has been realized 1-2 years after the initial shock, in line with the time pattern of house price effects.

Migration elasticity of house price growth. If we assume that net migration is the main transmission channel of house price growth between cities in the same migration network, we can compare the effect estimates for net migration and house price growth to get a rough estimate of the implied elasticity of house price growth with regard to population changes. Comparing the effect estimates for house prices and population, I obtain implied elasticities of house price growth with regard to population growth of about $0.23/0.05 \approx 4.6$ in the short-run (Column 6 of Table 4 and Column 3 of Table 5), and $0.63/0.187 \approx 3.4$ in the long-run (Column 1 and 4 of Table A3).

To put this elasticity in context, I compare this magnitude to the cross-sectional relationship between migration and house price growth during recent periods (see Appendix Figure A6). Comparing the change in average net migration rates and house price growth for each period to their values in 1991-1999, a 1 log point increase in the net migration rate was associated with increases in house price growth of 4.1 log points during the 2000-2007 boom, 2.1 log points during the bust 2008-2012, 1.9 log points during the 2012-2017 boom, and 2.8 log points during the 2000-2017 period as a whole. While these correlations are not directly comparable to my causal effect estimates, they suggest that the implied house price elasticity with regard to net migration of my spillover IV estimates is not far from historical relationships.

Reasons for moving. As additional qualitative evidence of why workers move between cities, I provide evidence from the Current Population Survey in Appendix H.3 on stated reasons for moving among migrants in the U.S. I find that housing reasons dominate among within-county migrants, while housing is stated with similar frequency as family or employment among moving reasons for inter-county migrants. Moreover, migrants moving for housing reasons are especially prevalent during the housing boom of the 2000s, and migration for housing reasons is an important driver of the variation in inter-county migration over time – in line with the migration spillover mechanism proposed in this paper.

[46]Note that migration flows are sufficient but not necessary for the spatial location choice mechanism to operate as suggested. In theory, if all cities were perfectly supply-constrained, any additional demand to migrate there would raise house prices until the excess demand by migrants was eliminated. That way, house price changes could propagate in equilibrium even with only modest or no migration flows being observed.

[47]See Column 4 of Appendix Table A3 for the point estimate.

4.2.2 Mortgage credit and construction effects

To highlight the key channels for how extensive margin increases in population could affect the housing market, this section estimates spillover effects on local mortgage lending and construction. I show that mortgage credit increases as a result of migration spillover shocks and that the dynamics of construction rationalize the house price increases.

Mortgage lending. To measure the effects on mortgage markets, I compute the number and dollar volume of purchase mortgages originated annually in each CZ in HMDA data. Columns 5 and 6 of Table 5 show the effects of the network house price shock on the log of housing purchase mortgages originated, as well as log housing purchase lending volume.

The results show that migration network spillovers significantly increase local lending activity, with a ten log point increase in network house prices leading to a 23 log point ($\approx 26\%$) increase in the number of purchase loans originated. Mortgage lending volume in USD (see Column 6) increases more than proportionally by 26 log points in response to a 10 log point shock, suggesting that average loan size is going up. In part this can be attributed to the fact that house prices are rising as a result of the network spillovers, but it is also consistent with other explanations, for example that migrating out-of-town buyers overpay as a result of being less well-informed (Chinco and Mayer, 2016).

Is this a plausible effect size? We can benchmark this effect by comparing it to the spillover effects on migration. In my HMDA sample, an average city and year see 1.1 mortgages originated per 100 residents. This means that the short-run increase in population of 0.5% that is estimated to result from a 10 log point network shock, multiplied by a U.S. homeownership rate of 67% in 2020, and divided by an average household size around 2.5,[48] should result in new loans corresponding to about $0.5 \times 0.67/2.5 = 0.13$ loans per 100 residents, or a 12% increase over existing loan volume.[49] Of course, lending to existing residents would also be expected to increase as a result of the impact of spillover shocks on house prices and within-city moving activity, and the total effect on the level of population in the long-run is four times larger than the short-run estimate used in this calculation. Taking into account these possible variations, the mortgage origination effect size appears plausible.

The dynamic effects on housing credit variables also behave in line with the migration and house price dynamics, with Panels (c) and (d) of Figure 6 showing the effect on new purchase mortgage originations and mortgage lending volume at different year horizons. The effects on the number of mortgages originated persist for several years but decline in magnitude to close to zero once population levels stabilize. The positive effect on lending volume also diminishes over time but remains significantly higher in the long run, consistent with the fact that house prices remain elevated in the long run as well.

Construction. Additional evidence on the changes in housing demand in response to network house price shocks comes from effects on the construction sector, measured by changes in construction permits issued. Column 4 of Table 5 and Panel (b) of Figure 6 show that spillover shocks increase construction, but only with a lag, as the positive effect only becomes significant

[48] This analysis assumes that the additional population has average characteristics and that all homeowners take out new mortgages. Average characteristics are from Census data at `https://www.census.gov/housing/hvs/files/currenthvspress.pdf` and `https://www.census.gov/data/tables/time-series/demo/families/households.html`.

[49] This may be an underestimate of the impact, as movers in the Current Population Survey for 1990-2017 are less likely to be married than non-movers (30% vs. 42%) and therefore might have a smaller family size.(Source: Author's calculation from CPS March samples obtained from IPUMS.)

one year after the initial shock.[50] The long-run effect on annual housing unit permits of a 10 log point network house price change is about 27 log points ($\approx 31\%$).

While this might seem large, it is important to put it in perspective on a per capita basis: over the 1990-2017 period, the Census reported an annual average of around 0.46 housing permits per 100 U.S. residents.[51] So, if net migration tends to increase the local population by $\sim$1.9 ppt in the long run as a result of a 10 log point network house price shock, and applying a household size of 2.5, then this implies a need for $1.9/2.5 = 0.76$ additional new units per 100 residents, or $0.76/0.46 = 165\%$ of average permit volume. Spreading this additional construction volume over 5 years, for example, the average city would have to increase annual construction by 33% to accommodate the increase in population – which is very close to the long-run construction effect that I have estimated.

The relatively large increases in construction required, relative to normal building volume, for even modestly sized population changes also rationalize why I find large house price effects of migration spillovers even in less geographically constrained cities: if a 1.9 ppt increase in population corresponds to a $> 33\%$ increase in construction volume over 5 years, it does not seem unreasonable for house prices to increase by 6.3 log points as the construction sector moves up the supply curve.

Housing vacancies during housing booms. The claim that congestion in construction markets plays an important role in understanding the propagation of house price increases across cities through migration contrasts with other narratives about how "overbuilding" was an important feature of the housing boom of the 2000s.[52] To provide additional evidence that congestion in construction markets plays a role in understanding the cross-section of housing booms, I therefore look at the relationship between vacancies and house price growth.

If migration congests housing markets, a given vacancy decline should be associated with a greater house price increase – and more so in more supply-constrained markets because they take longer to produce enough housing to accommodate the population increase. In contrast, under a narrative where overbuilding occurs during a boom and eventually attracts in-migration, vacancies should be positively associated with house price growth – and more so in constrained cities, as building is less responsive to a given house price increase.[53]

To test these predictions, I compute the change in vacancy rates and house price growth for housing boom periods relative to a baseline period, and compare more or less land-constrained cities.[54] The graphs of the relationship between vacancies and house price growth are shown in Figure 7 for owner-occupied housing (rental housing results are shown in Appendix Figure A8).

Vacancies are systematically lower in cities with higher house price growth in both periods -

[50]Note that this could either reflect a slow response by builders to rising house prices, or delays induced by the length of the permitting process.

[51]Calculated by dividing total units permitted 1990-2017 from `https://www.census.gov/construction/bps/`, by total U.S. population from `https://fred.stlouisfed.org/series/POPTOTUSA647NWDB`.

[52]As an example of this view, see e.g. Haughwout et al. (2012) (p. 70): "*While it is now clear that too much housing was built in the United States in the boom phase, identifying how much and where overbuilding occurred remain important issues...Our results suggest that 3 to 3.5 million excess housing units were produced during the boom. Excess housing production was a national phenomenon, but excess supply is positively related to housing supply elasticities*"

[53]On this topic, Haughwout et al. (2012) note that "*in markets with relatively elastic supply, bubbles should result in more new residential investment and consequently less of a price response.* "

[54]The baseline period is 1991-1997 and the boom periods are 2000-2007 and 2012-2017. This comparison can be computed for 98 CZs, which have data for all time periods. I group these cities into those with above- and below-median housing supply constraints based on their share of land unavailable for construction (Saiz, 2010; Lutz and Sand, 2019).

and this relationship is statistically significantly more negative in constrained cities. That is, the congestion effect of in-migration is worse if housing supply is harder to come by. This means that the evidence from vacancies is more consistent with migration causing house price growth through a shortage of housing supply than a speculative boom that leads to overbuilding which then entices migrants.[55]

Time-to-build delays. Another piece of evidence highlighting the significance of short-run congestion in construction markets is that time-to-build delays for residential housing construction increase during housing booms. This empirical pattern is shown in Figure 7. Here, I use data from the Census Survey of Construction to estimate average time-to-build from start to completion of construction projects completed in a given year for different regions. Note that all regions show increasing construction delays during the 2000s housing boom, falling delays during the Great Recession, and then increasing delays again during the recent 2012-2017 run-up in house prices.[56]

From the literature on investment cycles, we know that time-to-build delays for durable investment goods, such as dry bulk shipping vessels, can lead to large price cycles and predictable mean reversion in returns. The reason for this might be that individual firms do not take into account the endogenous investment behaviour of their competitors (Greenwood and Hanson, 2015). In the residential construction setting, this would correspond to developers responding to the initial housing market tightness documented above with eventual overbuilding at the end of the cycle – which would explain the narrative around the housing boom that focuses on the latter. The reason might be that investors neglect to account for the forthcoming additional supply of housing from other sources – either other builders, or the displacement of existing residents that I document in Section 4.2.3 below. At the same time, Kalouptsidi (2014) shows that procyclical time-to-build delays can exacerbate price volatility even in a rational model because investors cannot invest fast enough to restore equilibrium prices. Here, the existence of such time-to-build delays in housing construction provides an additional explanation for why housing supply does not immediately accommodate the migration spillovers on housing demand.

4.2.3 Educational displacement

This section provides evidence that economic booms that increase house prices can cause outflows of "displaced" workers. In particular, I focus on showing evidence of migration that results from college-educated workers moving into expensive cities and non-college educated workers moving out. Differences between education groups are salient because recent labor market trends have often benefited knowledge-intensive industries (Eckert et al., 2019).[57]

Displacement and housing supply constraints. First, I look at evidence that inflows of one education group can lead to outflows of the other group in cities with constrained housing markets. In Panel A of Figure 8, I plot gross out-migration rates of non-college workers (as %

[55]Note that I am not saying that Haughwout et al. (2012) and others are incorrect in remarking upon an increase in vacancies during the boom periods in the country as a whole. It simply turns out that, in the cross-section, the highest vacancy rates were not in the cities where the boom was happening, making it unlikely that speculative overbuilding was driving the vacancies - a fact that is obscured when looking at national averages. See Sumner and Erdmann (2020) for a similar argument.

[56]In a structural model calibration, Oh and Yoon (2020) show that these time-to-build increases at the height of the 2000s housing boom can be attributed to construction bottlenecks, whereas the continuation of delays at the beginning of the bust arises due to investor uncertainty about the housing market.

[57]Section 2.2 also provided additional support for group differences in the form of evidence that education groups may differ in their migration networks.

of local group population) over gross in-migration rates of college-educated workers by CZ for 2005-2017.[58] The graph shows a strong positive correlation between gross in- and outflows of these different education groups.[59]. If this correlation is in part mediated by housing markets, we should see differences in the degree to which college inflows result in non-college outflows in supply-constrained cities. In Panel B, I show only cities that fall into "high" or "low" constraint categories based on whether the local share of land available for construction (Saiz, 2010; Lutz and Sand, 2019) and regulatory constraints (Gyourko et al., 2019) are both above or both below median. The fitted line has a significantly steeper slope for highly constrained cities, consistent with larger house price increases leading to a greater displacement by inflows.[60]

Migration flows by education and housing cost. A corollary of this mechanism based on displacement is that workers in different education groups should on average be moving towards cities with different *levels* of housing costs. I test this by comparing house prices in origin and destination cities of migrants in different education groups. My measure of the expected change in quality-adjusted housing costs P^{Q}_{it} for migrants of each group s in year t is

$$\Delta_{Mig,s} P^{\mathrm{Q}}_{t} = \frac{\sum_k \sum_i \hat{m}^{ik}_{st} \left(P^{\mathrm{Q}}_{kt} - P^{\mathrm{Q}}_{it} \right)}{\sum_k \sum_i \hat{m}^{ik}_{st}},$$

where $\hat{m}^{ik}_{st}$ are imputed flows by education type from city i to city k.[61] This measure corresponds to the average difference in CZ house prices between destination and origin city across all migration flows for education group s.

The estimates of this term for the 2005-2017 period are shown in Panel C of Figure 8. As the time series graph shows, college-educated workers are consistently more likely to move towards more expensive cities than away from them during this period. In contrast, non-college workers on average move towards cities that are $\sim$ \$20K less expensive than where they are coming from at the peak of the housing boom. This displacement dynamic weakens during the recession and then picks up again in recent years. These data show that workers of different education levels are moving in opposite directions along the house price gradient, consistent with non-college workers being displaced from expensive cities towards cheaper locations as college workers move in.

[58] I use American Community Survey gross migration data, aggregated to commuting zones, by education, for 2005-2017. Annual migration data by education is only publicly available from the ACS for these years.

[59] This fact in itself is not novel, as it turns out that some cities experience consistently more population turnover than others (Coate and Mangum, 2019): gross in- and outflows are higher in high mobility cities, such as Las Vegas, NV, and Phoenix, AZ, than in low mobility cities like Providence, RI, or Philadelphia, PA. This leads to inflows being correlated with outflows in the cross-section of cities even within education groups, as Appendix Figure A9 confirms. This is also in line with other evidence, e.g. by Couture and Handbury (2017), that young college-educated workers have been more likely to move into some parts of urban downtown areas than non-college educated workers in recent decades. The contribution here consists of linking this displacement to housing markets.

[60] In Appendix Section H.5, I show that a similar dynamic can also be found when looking at the superstar cities case study: As a result of displacement, superstars experienced particularly rapid changes in their college share during the 2000-2007 boom, relative to other large cities.

[61] For this analysis, I follow Albouy and Ehrlich (2018) in constructing quality-adjusted housing cost indices P^{Q}_{it} for each city, by adjusting raw house prices using hedonic regressions on housing characteristics – see details in Appendix F.7. Moreover, Appendix F.6 shows how I construct imputed migration outflow shares μ^{ik}_{st} by education group s from city i to city k, for the 2005-2017 period, using ACS data. Imputed total flows by education type are then computed as $\hat{m}^{ik}_{st} = \mu^{ik}_{st} * L_{is,t-1}$, where $L_{is,t-1}$ are workers older than 21 years in the ACS data of education group s, and in city i in the previous period.

4.3 Importance of migration spillovers for house prices

While the migration spillover mechanism is general, in the sense that *any* shock to an origin city's housing markets can propagate through migration flows to other cities, this section considers the effects of two particular shocks through this spillover channel. I show that shocks to other cities in the form of interest rate effects on house prices, as well as from exogenous wage changes, induce sizeable spillover effects on other cities that can be quantified using my estimates. Moreover, I show that migration links are useful for predicting co-movement in house prices among cities in general. However, the reduced-form estimates do not allow for studying the aggregate effect of changes in mobility on house price dynamics – that question will instead be addressed by the structural model estimation.

4.3.1 Quantifying spillovers from "easy credit" and wage shocks

To get an estimate of how important spillover effects are for understanding the effect of historical shocks on house price growth patterns, I estimate the implied spillovers from well-known house price shocks in the literature. This quantification also illustrates how the migration spillover methodology can be applied more generally to understand the spatial externalities from localized shocks. I focus on shocks to house prices in the run-up to the housing boom of the 2000s that consist of (1) easy credit in the form of low real interest rates, and (2) industry wage shocks.

The effect of decreases in the user cost of housing as a result of lower real interest rates is frequently cited as one of the causes of increases in housing demand in the late 1990s and early 2000s, although the exact size of the effect is subject to debate (Himmelberg et al., 2005; Glaeser et al., 2012). Here, the focus is on the *differences* across cities in the effect of interest rate changes on house prices that result from differences in land constraints interacting with housing demand changes. I follow Glaeser et al. (2012) in estimating the size of these effects.[62] Similarly, many researchers have identified differences in city exposure to industry wage trends interacting with supply constraints as an important driver of differential house price changes.[63] Here, I show the extent to which the spillovers from these shocks can explain part of the house price dynamics in *other* cities.

For each shock, I first estimate a "first stage" consisting of the effect of that shock on house prices in other cities that are part of a focal city i's migration network, controlling for the baseline controls that are also applied in my spillover regressions, as well as any direct effects of the shock on city i. Then, I apply the baseline short-run spillover effect estimate $\widehat{\tilde{\eta}^{nw}}$ (see Column 6 of Table 4 to the predicted exogenous 1995-2007 change in network house prices for each city i. The result is an estimated spillover effect on city i for 1995-2007 from each shock to other cities, which I then compare to the actual house price growth and house price beta of city i over that period.[64]

The ability of these predicted spillover effects to explain actual CZ house price growth in 1995-2007 is shown in Figure 9. Panels (a) and (b) show that the cross-sectional variation in spillovers from interest rate effects and wage shocks in other cities can explain 32% and 12%, respectively,

[62] See Appendix Section F.3 for details on constructing the interest rate spillover shocks.

[63] See, e.g. Saiz (2010); Paciorek (2013); Diamond (2016).

[64] As the wage shocks were also used in the spillover estimation, their construction closely follows the IV methodology discussed earlier: Table 3 shows the first-stage estimates used to predict the effect of the (residualized) instruments $NWP_{it}^{B,resid}$ and $NWP_{it}^{Bx,resid}$ on network house price growth, and the direct effect controls for city i consist of the wage shocks to city i and their interaction with land constraints that are already included in the set of baseline controls.

of the variation in CZ house price growth over this period. Moreover, the predictive linear fit for the spillover from each of the shocks has a positive and significant slope.[65] Of course, particular shocks to other cities can only ever partially explain the house price growth in city i that is caused by *all* shocks – observed or unobserved – together.

To understand whether these spillovers contribute positively or negatively to the house price beta in affected cities, I compute national trend betas for the 1995-2007 period for each city. The effect of the spillovers on these betas is shown in Panels (c) and (d). Again, the effect of the predicted spillover shocks on the cross-section of betas is positive and significant for both interest rate and wage shock spillovers.[66] A simple regression of betas on the residualized spillover shocks shows an R-squared of 32% and 14% for the interest rate and wage shock spillovers. Of course, this analysis could be extended to other shocks that play a part in determining house prices over this time period, providing a tool for other researchers interested in exploring spatial spillovers arising from a shock of interest.

4.3.2 Migration links as a predictor of house price correlations

This section confirms the importance of spillovers through migration links more generally by showing that migration links are good predictors of which cities experience correlated house price movements, and that they have predictive power for house price correlation that is not explained by other links, such as similarities in industry structure.

Sinai and Souleles (2013) first noted that pre-Great Recession U.S. migration tended to occur between MSAs that have correlated house price cycles. This correlation can be explained by the spillover mechanism that I propose in this paper: If migrant flows through migration networks transmit local house price dynamics, we would expect house prices to be correlated more between cities that have stronger migration links. I expand on the analysis of Sinai and Souleles (2013) by showing that migration links are in fact a *better* predictor of house price correlations than other plausible measures of inter-city links.[67]

As a first step, I compute the pairwise correlation of house price growth between CZs over 1990-2017, using the log change in the FHFA repeat-sales index as my measure of house price growth. Then, I compute for each CZ i the *weighted* average correlation with all other CZs $k \in N$, which is given by

$$E_i^m(corr(\Delta \ln P_i, \Delta \ln P_k)) = \sum_{k=1}^{N} w_{ik}^m \cdot corr(\Delta \ln P_i, \Delta \ln P_k).$$

The weights w_{ik}^m represent a particular measure m of inter-city links considered. If the inter-city link measure m better predicts pairwise house price correlations, the link-m-weighted expected correlation should be larger. The weights are always normalized to sum to one.

My main measure of links reflects migration connections: the average share of all migrants out of city i in the IRS migration flow data that move to city k. I compare the ability of migration

[65] Note that the predicted spillover effects are residualized with regard to direct effects of the shocks and all baseline controls, which means that the absolute magnitude of these spillovers cannot be straightforwardly interpreted.

[66] Note that it is not automatically the case that a positive house price growth effect coincides with a positive beta effect: If shocks to a city's migration network are positive on average but timed in a way that correlated negatively with national house price growth, the house price growth and beta relationships can have slopes of opposite signs.

[67] In addition, I also expand on the analysis by Sinai and Souleles (2013) by including the decade since the Great Recession in the time frame, and expand the sample geographically by including commuting zones that cover the entire continental U.S.

weights to predict house price correlations to a number of other competing measures of inter-city links: (1) An inverse-distance weighted measure that reflects geographic proximity; (2) A social connectedness index (SCI) based on Facebook friendship links between geographic areas that reflects the strength of social networks between cities (see Bailey et al. (2018b)); (3) A destination population-weighted measure of house price correlations that reflects the outsized role of large cities in the U.S. economy and therefore perhaps also in housing cycles; (4) A measure of the similarity in industry structure between cities;[68] (5) An equal-weighted measure that simply reflects a city's average correlation with other cities' house prices; (6) A measure of trade flow links between cities that represents industry linkages and the propagation of economic shocks through input-output networks.[69]

The expected house price correlation for each CZ under these different weights is shown in Figure 10. The graph sorts CZs by their migration-weighted house price correlations (which therefore form a line), and then plots for each CZ i the expected house price correlation with all other CZs under the alternative weights described above. A higher value indicates that the link captured by those weights is more strongly associated with a co-movement in house prices.

As is clear from the graph, the correlation in house prices between cities that have stronger migration links is almost everywhere larger than the house price correlation predicted by distance, social network links, city populations, industry similarity, trade flows or equal-weighting all other cities. This visual result is confirmed by the average expected correlations for each link shown in the legend of the graph. This predictive outperformance suggests migration links can *not* simply be considered as proxies for one of these other city connections.[70] For example, even though migration links are likely to be at least in part correlated with economic linkages between cities – which are proxied here by the relative value of trade flows – the fact that migration links are a better predictor of house price correlations than trade links suggests that the explanatory power of migration links cannot be reduced to being merely a reflection of a trade channel.

To quantify the predictive importance of migration links, in Appendix Section H.2, I run horserace regressions between these different link measures and find that migration links have strong and significant predictive power for house price correlations even when accounting for all other measures together.[71] The implied effect of stronger migration links (controlling for all other links, as well as the overall tendency of cities to send or receive migrants) is that an 10 ppt increase in the share of migrants from city i that go to city k increases the correlation in house prices with that city by 6.7 ppt. Note that this is very similar to the correlation in house price growth implied by the causal spillover effect estimate: in the migration network regressions, if the migration exposure to a city increased from zero to 10 ppt, then the long-run spillover elasticity of 0.63 (see Column 1 in Table A3) suggests that there will be a 6.3 ppt long-run correlation with house price shocks in the upweighted city due to spillover effects. While these analyses are not directly comparable due to the different time horizon, approach, and migration weights, it is encouraging that this predictive analysis yields a similar order of magnitude for the importance of migration links. Moreover, the shifts in migration shares considered here are necessarily partial

[68] More details on these measures can be found in Appendix Section H.2.

[69] See Appendix Section F.9 for details on the construction of this measure, which reflects the share of a city's domestic value of outgoing trade flows that is going to another city.

[70] In Appendix Figure A7, I repeat the same analysis omitting any CZs within a 50 or 150 miles distance of one another from the calculation of the averages to capture only the long-distance effect via these links - and find very similar results.

[71] The only measure that is not directly included in the horseraces is the trade flow link measure, which is only available for a small subset of CZ pairs, representing $< 10\%$ of the full sample.

equilibrium changes, as they do not account for second-order effects on other cities. In order to assess the full general equilibrium impact of changes in mobility, I will make use of the structural dynamic spatial equilibrium model discussed in the following sections.

5 Dynamic Spatial Equilibrium Model

5.1 Model overview

The previous sections showed the reduced-form causal spillover effect of cities on one another. Those effect estimates represent average effects over the time period (1990-2017), taking the effect of structural parameters, such as house price elasticities and migration costs as given. However, migration costs have likely changed over time, contributing to a secular decline in domestic U.S. migration rates (Molloy and Smith, 2019). Similarly, political decisions that increase housing supply constraints, such as zoning rules, seem to have led to lower elasticities of housing supply over time (Ganong and Shoag, 2017; Gyourko et al., 2019), engendering a recent backlash in the form of the "yes in my backyard" (YIMBY) movement which calls for reductions in housing supply constraints to improve housing affordability (Dougherty, 2020).

At the same time, the reduced-form estimates do not capture any higher-order effects from spillovers: when migrants cause an increase in house prices at their destination, this may in turn cause some local residents to move on, leading to migration into another set of cities, and so on. As a result, the effects of migration on the aggregate pattern of house price growth responses to an economic shock depend on the full network of city migration links, including higher-order connections. To understand what role changes in migration costs and house price elasticities play in contributing to the aggregate distribution of house price growth, – and to understand the effect of the aforementioned trends – I need to model the equilibrium mechanism that transmits shocks between cities. This mechanism includes workers' endogenous location choices, which are the driving force for migration flows that respond to house price increases.

I therefore develop a quantitative dynamic spatial equilibrium model that captures the key mechanisms for understanding house price spillovers effects: On the one hand, the model builds on standard location choice models with heterogeneous worker groups (e.g. Diamond (2016)), but incorporates methods from the dynamic discrete choice literature to allow for forward-looking migration choices, as workers might anticipate future trends in city characteristics. Similar methods have recently been applied in a location choice context by Diamond et al. (2017), Coate and Mangum (2019), and Almagro and Domínguez-Iino (2020), among others. In addition, I allow for a rich structure of different bilateral migration costs between cities that shapes differences in cities' migration networks and exposure to other cities' shocks, using insights from Caliendo et al. (2019) to show how to estimate this model without knowing the level of moving costs.

On the other hand, the model incorporates heterogeneous housing supply constraints. As cities' housing markets will differ in the degree to which changes in housing demand drive up house prices, this is an important part of understanding differences in house price dynamics across cities. As far as I know, this is the first application of a dynamic location choice model with heterogeneous housing supply to inter-city migration in the U.S.

To simulate what would happen under different scenarios for migration costs and supply constraints, I build on recent innovations in estimating counterfactuals in dynamic location choice models. Ahlfeldt et al. (2020) noted that in dynamic spatial models with migration costs observed allocations may not be stationary as long-run equilibrium is not immediately attained. Nonetheless, they show that comparative statics can be computed by estimating and comparing stationary equilibria implied by baseline observables and counterfactuals. I adapt this idea to my model with repeated location choices and flexible forward-looking expectations, showing that one can use the "dynamic hat" methods of Caliendo et al. (2019) to solve for stationary equilibria. The counterfactual transition paths to these stationary equilibria can then be compared to understand the equilibrium impact of changes in parameters.

Along the way, I develop a methodology for working with a limited sample of migration data. In the U.S. context, origin-destination data on migration flows is available from multiple data sources that have different strengths and weaknesses. I use an "empirical Bayes" approach to impute smoothed migration probabilities from censored flow data. I reduce the noise of migration flows by education group by using machine learning and additional information about flow totals to generate predicted bilateral flows. These predicted flows are then combined in a "Bayes shrinkage" approach with the raw data to reduce the noise that comes from a limited sample size.

To be able to extend counterfactuals to past time periods where detailed flows by education group are not observed, I also develop a method to use the model structure together with observable city characteristics and total migration information to simulate the model backwards and impute unobserved flows by education groups that are consistent with the observed data.

This section will present the dynamic spatial equilibrium model. The following section then discusses how to estimate and calibrate the key model parameters. After that, I will discuss the parameter estimates, and proceed to explain the methodology for simulating counterfactuals in this model before considering the scenarios of interest, in which I show the effect of different migration costs and changes in housing supply constraints on house price dynamics.

5.2 Within-period worker preferences

Flow utility. All workers live in some city $i \in N$ at the beginning of period t. They have log utility over a Cobb-Douglas aggregator function of tradable consumption goods with uniform prices across all locations, and local nontradable housing, with per-period unit cost Q_{it}. Workers belong to different skill groups $s \in \mathcal{S}$, which differ in their preferences for amenities and the expenditure share of nontradables α_s, both of which have been found to affect group-specific location choices (Bayer et al., 2016; Diamond, 2016; Almagro and Domínguez-Iino, 2020). The price of these nontradables is assumed to vary with housing costs.

The indirect flow utility for a worker ω in location i and of type s can therefore be written as:

$$U_{ist}(\omega) = \underbrace{\ln A_{ist}}_{\text{Amenities}} + \underbrace{\ln \widetilde{W}_{ist}}_{\text{Work utility}} - \alpha_s \underbrace{\ln Q_{it}}_{\text{Housing cost}} \tag{8}$$

. Here, $\widetilde{W}_{ist}$ is the expected local wage index and A_{it} captures residential amenities from living in i.

Heterogeneity. To model the heterogeneity among workers in their attachment to different locations, I assume that in each period workers draw an idiosyncratic location amenity shock $z_{ist}(\omega)$ for each location i, which is Type 1 extreme value (Gumbel) distributed with variance $\frac{\pi^2}{6}$, which is scaled by parameter θ_s. A smaller θ_s corresponds to less heterogeneity in idiosyncratic preferences among workers and therefore a higher sensitivity to differences in common factors between locations. The idiosyncratic location amenity is realized right after the moving decision.

In addition, whenever a worker moves between two locations i and k, she has to pay an additive and time-invariant moving cost τ_s^{ik} in utility units in the same period. This moving cost can depend on the skill group s, which means that migration networks can vary by group, in line with the findings shown in Section 2.2 that college and non-college educated workers differ empirically in the cost attached to overcoming differences in characteristics between origin and destination city. This moving cost generates geographic differences between cities and is important for explaining why the geographic distribution of shocks matters.

Industry choice. After choosing a location, each period workers also draw an idiosyncratic industry preference shock $z_{i\iota st}(\omega)$ and solve an industry choice problem of the form

$$\max_{\iota} \left[W_{i\iota t} \cdot z_{i\iota st}\right],$$

where $W_{i\iota t}$ are wages in industry $\iota \in \mathcal{I}$ in location i, and $z_{i\iota st}$ is an industry preference shock. I assume that this idiosyncratic industry preference shock is Fréchet-distributed with CDF

$$F(z_{i\iota st}) = e^{-\epsilon_{s\iota i} z_{i\iota st}^{-a}},$$

where the scale parameter $\epsilon_{s\iota i}$ determines the average suitability of group s workers in industry ι and location i, and the parameter a captures the inverse heterogeneity in industry preferences. The means $\epsilon_{s\iota i}$ of these shock distributions vary by location and group but not over time. This parameter reflects skill requirements in different industries: workers who have more education may have to expend less effort to work in particular industries, and the same industry in different locations may skew more or less towards favouring more skilled workers in terms of its work conditions.

Conditional on having chosen location i, the probability that a group s worker chooses to work in industry ι is therefore given by

$$\pi_{is\iota t} = \frac{W_{i\iota t}^{a} \epsilon_{s\iota i}}{\sum_{\iota=1}^{\mathcal{I}} W_{i\iota t}^{a} \epsilon_{s\iota i}}. \tag{9}$$

Importantly, location choices are made before the realizations of idiosyncratic productivity shocks are drawn. Therefore, the relevant variable for group s location choices is the expected income index $\widetilde{W}_{ist}$, which by the properties of the Fréchet distribution is given by

$$\begin{aligned} \widetilde{W}_{ist} &= E_z[\max_{\iota} [W_{i\iota t} \cdot z_{i\iota st\omega}]]. \\ &= \Gamma\left(\frac{a-1}{a}\right)\left(\sum_{\iota=1}^{\mathcal{I}} W_{i\iota t}^{a} \epsilon_{s\iota i}\right)^{\frac{1}{a}}, \end{aligned} \tag{10}$$

where $\Gamma(\cdot)$ is the Gamma function. Solving for changes in this work utility index over time, it evolves according to[72]

$$\Delta \ln \widetilde{W}_{ist} = \frac{1}{a} \ln\left(\sum_{\iota=1}^{\mathcal{I}} \pi_{is\iota t} \frac{W_{i\iota t}^{a}}{W_{i\iota,t-1}^{a}}\right). \tag{11}$$

[72]Note that this model could also be easily adapted to accommodate the possibility of unemployment as an alternative "industry" that a worker might end up in.

Model timing. To summarize, the timing in the model is illustrated in Figure 11: In each period, a worker in location i receives the idiosyncratic location shock draws, chooses a destination city k – which may be the same as the origin city – and executes the move. Next, the worker receives an industry preference shock in the chosen location and picks an industry to work in. Once that decision has been made, the worker pays moving costs, realizes the preference shocks, and then works and consumes in the chosen location until the end of the period.

5.3 Forward-looking location choice

An important consideration in using annual data rather than decadal frequencies is that workers' time horizons in evaluating which city to move to will likely also take into account future periods' city characteristics. Therefore, I model moving decisions as forward-looking: Workers anticipate their ability to make migration decisions in future periods as well as the future flow utility obtained from different locations.

That is, workers consider the conditional value function υ_t^{is} of each location i when making their moving decisions rather than only the flow utility. They solve the following problem:

$$\max_{k\in N}\left\{\overbrace{\underbrace{U_{kst}}_{\text{Flow utility}}+\underbrace{\beta E\left[V_{t+1}^{ks}\right]}_{\text{Option value}}}^{\text{Conditional utility } \upsilon_t^{is}}-\underbrace{\tau^{ik_s}}_{\text{Moving cost}}+\underbrace{\theta_s z_{ks}(\omega)}_{\substack{\text{Idiosyncr. location}\\ \text{preference: T1EV}}}\right\}.$$

The conditional utility υ_t^{is} of location i for a group s worker after the location decision has been made but before industries have been chosen is defined as

$$\upsilon_t^{is} = U_{ist} + \beta E[V_{t+1}^{is}], \tag{12}$$

where U_{ist} is the flow utility defined in Equation 8, and the second term in this expression corresponds to the expected ex-ante value function for the next period. The ex-ante value function is defined as the expected value, at the beginning of the period, of being in location i before the idiosyncratic shocks are drawn:

$$V_t^{is} = E_z\left[\max_{\{k\}_{k=1}^N}\left\{\upsilon_t^{ks} - \tau_s^{ik} + \theta_s z_{kst}\right\}\right],$$

where the outer expectation $E_z[\cdot]$ is over the distribution of z_{kst}.

Using the properties of the Gumbel distribution (see Hotz and Miller (1993)), we can rewrite this expressions as

$$V_t^{is} = \theta_s\gamma + \theta_s \ln\left(\sum_k^N \exp(\upsilon_t^{ks} - \tau_s^{ik})^{\frac{1}{\theta_s}}\right) \tag{13}$$

where γ is Euler's constant. The probability of deciding to relocate to any city k from city i can be written in the standard logit form as

$$\mu_{st}^{ik} = \frac{\exp(\upsilon_t^{ks} - \tau_s^{ik})^{\frac{1}{\theta_s}}}{\sum_j^N \exp(\upsilon_t^{js} - \tau_s^{ij})^{\frac{1}{\theta_s}}}. \tag{14}$$

To express the location option value in terms of observables, we can take logs of the conditional moving probability μ_{st}^{ik} for a particular destination choice k and rearrange to obtain

$$\theta_s \ln\left(\sum_k^N \exp(\upsilon_t^{ks} - \tau_s^{ik})^{\frac{1}{\theta_s}}\right) = \upsilon_t^{ks} - \tau_s^{ik} - \theta_s \ln \mu_{st}^{ik}.$$

Then, substituting this expression into the ex-ante value function in equation 13, and using the definition of the conditional value function, I obtain

$$\begin{aligned} V_t^{is} &= \upsilon_t^{ks} - \tau_s^{ik} + \theta_s\gamma - \theta_s \ln \mu_{st}^{ik} \\ &= U_{kst} + \beta E[V_{t+1}^{ks}] - \tau_s^{ik} + \underbrace{\theta_s\gamma - \theta_s \ln \mu_{st}^{ik}}_{\substack{\text{Non-optimal choice} \\ \text{adjustment}}} \end{aligned} \tag{15}$$

That is, the ex-ante value of having a choice of where to move to from city i can be written as the utility obtained from a particular choice k (the first three terms), adjusted by the degree to which k is suboptimal, which will be reflected in the actual probability of workers choosing k. This result follows directly from the properties of the Gumbel preference shock distribution (Arcidiacono and Miller, 2011).

This rewriting of the ex-ante value function in the dynamic discrete choice problem into the sum of the certain value of a particular choice, and an adjustment term that is a function of observable choice probabilities, has been used widely in the trade literature (e.g. Artuç et al. (2010); Caliendo et al. (2019); Traiberman (2019)). Note that our choice of destination k was arbitrary for this derivation. In the estimation below, I will use this fact to write the option value of a location as a function of the value of choosing any *arbitrary* destination choice k, including the origin location, adjusted by a function of the empirical migration share for that destination. This is useful, because migration shares in future periods are empirically observable, even in finite data.

5.4 Production and labor demand

Production technology. Firms in each industry ι are indifferent between workers of any group s and produce output $Y_{\iota t}$ that consists of tradable goods differentiated by industry. The production technology uses local labor as the only input and has constant returns to scale:

$$Y_{\iota t} = X_{\iota t} L_{\iota t},$$

where $L_{\iota t} = \sum_s^{\mathcal{S}} L_{\iota ts}$ is the sum of local industry ι employment across workers in groups $s \in \mathcal{S}$. Here, $X_{\iota t}$ denotes the local productivity of firms producing goods in industry ι. Moreover, I assume that there is perfect competition in input markets such that workers earn their marginal product.

Tradable goods demand. I model the local demand for differentiated tradable goods following Armington (1969).[73] I assume that, within the tradables category, the consumption utility in location i over which different groups s have Cobb-Douglas preferences is an aggregator of local

[73] This demand structure is common in the trade and economic geography literature – see, e.g. Arkolakis et al. (2012), or Adao et al. (2019) for a similar recent application.

industry good consumption of the form

$$C_{it} = \prod_{\iota=1}^{\mathcal{I}} C_{i\iota t}^{\gamma_\iota}, \tag{16}$$

where $C_{i\iota t}$ – the industry ι goods consumption in city i – is in turn an aggregator of the utility obtained from the consumption of differentiated goods from other locations j given by

$$C_{i\iota t} = \left(\sum_{j=1}^{N} (c_{ij\iota})^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}},$$

where $\sigma \in (1, \infty)$ and $c_{ij\iota}$ is the consumption of industry ι goods from city j in city i. Trade is assumed to be costless and product markets are competitive, so goods from a particular location cost the same in all other locations.

Labor demand. We can also take into account the possibility of agglomeration effects from greater employment, which have been previously documented in the literature (Diamond, 2016). Assume that

$$X_{i\iota t} = \widetilde{X}_{i\iota t} L_{i\iota t}^{\alpha},$$

so productivity in each industry changes with the local employment level, and $\widetilde{X}_{i\iota t}$ is the residual local industry productivity. Then, local industry wage changes become

$$\Delta \ln W_{i\iota t} = \eta^{LD} \Delta \ln L_{i\iota t} + \frac{1}{\sigma} \Delta \ln \tilde{D}_{i\iota t}, \tag{17}$$

where $\widetilde{D}_{i\iota t} = \left(\widetilde{X}_{i\iota t} P_{\iota t} \right)^{\sigma-1} \gamma_\iota$ summarizes local industry demand and productivity shocks, and $\eta^{LD} = \left(\frac{\alpha(\sigma-1)-1}{\sigma} \right)$. Here, $P_{\iota t}$ is the composite price index for goods from industry ι. See Appendix Section E.3.1 for details on the derivation.

Equation 17 is the labor demand equation that I will take to the data. As regards timing, I assume that wages are determined at the beginning of the period, before migrants make their location choices.

5.5 Housing markets

Housing valuation. An important limitation in U.S. housing data is the lack of a comprehensive panel of rents at a city level at sub-decadal frequencies. However, in practice the user cost of housing for households that are not homeowners depends on the level of annual rents, rather than the level of house prices. Substituting changes in the latter for changes in the former implicitly assumes that the valuation of houses is myopic and does not take into account other factors, such as the expected future growth of rents. See Appendix Section E.3.2 for a discussion on what different informational assumptions imply for the relationship between rents and house prices.

Nonetheless, the myopic pricing model is an important benchmark case because it corresponds closely to the implicit assumption in most static spatial equilibrium models. For instance, in constructing time series of changes in housing costs, rent data are often imputed as a constant multiple of house prices (see, e.g. Diamond (2016); Ganong and Shoag (2017)), or it is assumed that rents can be mapped directly into contemporaneous house prices using interest rates (e.g.

Piyapromdee et al. (2014)). Similarly, spatial models in economic geography and trade often employ models where prices for real estate clear the market each period but are not forward-looking, such that per-period housing cost and house prices are the same (e.g. Ahlfeldt et al. (2015); Caliendo et al. (2019); Sturm et al. (2020)).

Due to the data limitations discussed above, I unfortunately cannot assume fully rational price growth expectations that are jointly varying across time and space. Instead, I assume that valuations can vary over time nationally by a common factor α_t, and can vary by a fixed factor α_i across cities – but do not allow for differences in local valuation trends as these cannot be identified separately from the city-level effects that I am interested in. That is, I assume that

$$P_{it} = Q_{it} e^{\alpha_t} e^{\alpha_i}. \tag{18}$$

Housing demand. The worker's preferences over housing imply that housing expenditure is a constant share α_s of the wage bill for group s. Therefore, we can write real housing demand as

$$H_{it}^D = \frac{e^{\alpha_t} e^{\alpha_i}}{P_{it}} \left(\sum_{\iota} \sum_{s} \alpha_s W_{\iota i t} L_{is\iota,t} \right),$$

where H_{it}^D is the real quantity of housing demanded and $L_{is\iota t}$, is employment of group s in industry ι.

Housing supply. Housing developers are assumed to be price-takers who sell housing consumption units at marginal production cost. Individual developers are marginal with regard to competition, but their output as a whole can be written in the form of a representative firm. The local housing supply H_{it}^S is assumed to be provided using a Cobb-Douglas technology that combines perfectly mobile construction capital with local land and local construction productivity. Differences in the local housing supply elasticity are generated by the assumption that the cost of the local land input increases with the size of the total housing stock with a location-specific elasticity. Solving the developer's cost minimization problem (see Appendix Section E.3.3 for details), the local inverse housing supply function can be written as

$$P_{it} = \phi_i \phi_t (H_{it}^S)^{\phi_i^H}, \tag{19}$$

where the variables other than the housing quantity supplied are functions of model parameters that only vary along the dimension indicated in their index.

Housing market equilibrium. Setting total housing demand equal to the housing stock, i.e. $H_{it}^D = H_{it}^S$, the housing market clears at a price

$$P_{it} = \tilde{\phi}_i \tilde{\phi}_{it} \left(\sum_{\iota}^{\mathcal{I}} \sum_{s}^{\mathcal{S}} \alpha_s W_{\iota i,t} L_{is,\iota,t} \right)^{\tilde{\phi}_i^H}, \tag{20}$$

where the responsiveness $\tilde{\phi}_i^H$ of house prices with regard to housing expenditures is higher if the city is more land-constrained.[74]

[74]The parameters here are given by $\tilde{\phi}_i^H = \frac{\phi_i^H}{1+\phi_i^H}$, and $\tilde{\phi}_i = \kappa_i^{\frac{-1}{(\alpha_K+\alpha_A)(1+\phi_i^H)}}$, and $\tilde{\phi}_{it} = \left((r_t^K)^{\alpha_K} \alpha_A^{-\alpha_A} \alpha_K^{-\alpha_K}\right)^{\frac{1}{(\alpha_K+\alpha_A)(1+\phi_i^H)}} e^{\frac{(\alpha_t+\alpha_i)\phi_i^H}{1+\phi_i^H}}$.

5.6 Transition dynamics

To summarize the model dynamics, this section characterizes what this model implies for the transition dynamics of the economy between periods.

Given the current period employment distribution for each group across locations and industries, migration costs, wages, housing costs and house prices, location amenities, and local industry demand and productivity fundamentals, as well as national house price valuation trends and expectations of location option values, the economy evolves as follows: next-period employment levels, wages, and housing costs across all locations $i \in N$, industries $\iota \in \mathcal{I}$ and skill groups $s \in \mathcal{S}$ are determined by

1. Workers choosing their location optimally when given the chance to migrate, such that flows between cities are given by Equation 14, which depends on each location's flow utility and option value as defined in Equations 8 and 15. Then, employment in each location and group evolves according to

$$L_{i,t+1} = \sum_s L_{is,t+1} = \sum_s \sum_k \mu_{st}^{ki} L_{kst}.$$

2. Local industry-level wage changes are determined by the inverse labor demand in Equation 17, and the ex ante work utility index for each location and education group updates based on Equation 11.

3. House prices are determined by the housing market equilibrium condition in Equation 20, and house prices are linked to rent growth through the valuation function in Equation 18.

The structural parameters of this model consist of the vector $\{\theta_s^{-1}, \eta^{LD}, \tilde{\phi}_H^i, \alpha_s, \beta, a\} \quad \forall (s, i)$. In the next section, I show how I calibrate or estimate these parameters to simulate the model.

6 Quantitative Estimation Approach

This section describes how to take the dynamic model described above to the data. First, I explain how I calibrate auxiliary parameters based on the literature. Second, I describe the approach to estimating the main structural parameters of interest . Last, I describe how I clean and construct the data with which to estimate the model.

6.1 Calibrated parameters

Where possible, I estimate the key parameters of the structural model. Before I proceed to explain how those parameters are identified, this section describes the values chosen for the auxiliary parameters that are calibrated ex ante. An overview of the values of all key structural parameters is shown in Table 6.

Discount factor. As the discount factor β is difficult to identify in dynamic discrete choice models (Manski, 1993; Rust et al., 1994) and I do not have credible instruments to separately identify variation in current and future values (Magnac and Thesmar, 2002; De Groote and Verboven, 2019), I follow the applied dynamic discrete choice literature (e.g. Traiberman (2019)) and set the discount factor exogenously. In particular, I follow Kennan and Walker (2011), Dix-Carneiro

(2014), and Artuc et al. (2020) in setting $\beta = 0.95$ as the annual discount factor. However, none of the main results are very sensitive to this choice.

Housing expenditure shares. I calibrate the nontradable expenditure shares α_s of local nontradables associated with housing costs based on existing results in the literature. I use the estimates by Diamond (2016) and set $\alpha_{col} = 0.63$ and $\alpha_{nc} = 0.68$ for college and non-college workers.[75] In one variation, I also assume common expenditure shares of 0.65. In order to impose the model implication that the combined flow utility enters with a common elasticity into the location choice equation, I combine the log wage utility growth and housing cost growth into a log real wage growth term

$$\Delta \ln \mathcal{R}^w_{ist} = \Delta \ln \widetilde{W}_{ist} - \alpha_s \ln Q_{it}, \tag{21}$$

using these calibrated expenditure shares.

6.2 Renewal actions

In estimating the location choice parameters θ_s, I have to address the fact that the conditional value function in Equation 12 contains future location value terms $E[V^{is}_{t+1}]$. The intuition for the presence of these terms is that workers might move because a city will *become* expensive or high wage or high in amenities etc. in the future. However, with finite data, we face the issue that we do not fully observe future value functions V^{is}_{t+1}.

In order to identify the parameters θ_s without observing all state variables, I therefore follow an approach based on Scott (2013) that has recently found application in research on urban location choices (Diamond et al., 2017; Almagro and Domínguez-Iino, 2020). It is based on using future realized values of observables as a proxy for agents' expectations that influence current decisions. It also assumes finite dependence in actions by modelling location choices as "renewal actions" where the same choices leave agents in the same state with regard to future value functions, which will allow us to difference out future value terms.

Renewal action illustration. The intuition for a one-period ahead renewal action in a migration setting is illustrated in Figure 12. Workers 1 and 2 both start out in city i at the beginning of period t, but Worker 2 chooses to move to city k in period t, while Worker 1 remains in place and only moves to city k in period $t + 1$. Migration in period $t + 1$ to city k leaves Worker 1 in the same state as Worker 2, who moved earlier. The renewal action assumption is that, going forward, both workers have the same expected value of $t + 2$ actions and outcomes, given by $E[V^{ks}_{t+2}]$. However, their utility in period t may vary along the different paths taken, as they experience wages and rents in different cities during period t, and pay moving costs at different times. Therefore, the differences in city characteristics at time t, together with the choice probabilities for the different paths, can be used to estimate the sensitivity of migration choices to city observables.

To derive an estimating equation based on this renewal state approach, assume that two households of type s, both in location i at the beginning of period t, make some location choices k and l. We can take the difference between the logs of the conditional probabilities of these choices to

[75] The same calibration choice is made in a recent paper by Fajgelbaum and Gaubert (2020). These expenditure shares are also very close to the share of consumption expenditures with prices that are affected by house prices computed in Moretti (2013).

obtain

$$\begin{aligned}\theta_s \ln\left(\frac{\mu_{st}^{ik}}{\mu_{st}^{il}}\right) &= (\upsilon_t^{ks} - \upsilon_t^{ls}) - (\tau_s^{ik} - \tau_s^{il}) \\ &= (U_{kst} - U_{lst}) + \beta\left(E[V_{t+1}^{ks}] - E[V_{t+1}^{ls}]\right) - (\tau_s^{ik} - \tau_s^{il}) \\ &= (U_{kst} - U_{lst}) + +\beta\left(V_{t+1}^{ks} - V_{t+1}^{ls}\right) - (\tau_s^{ik} - \tau_s^{il}) + (\xi_t^{ks} - \xi_t^{ls}) \qquad (22)\end{aligned}$$

where $\xi_t^{ks} = \beta(E[V_{t+1}^{ks}] - V_{t+1}^{ks})$. That is

Expectational error assumption. In the last line of Equation 22, I replaced the expectation operator with the realized value of the future value function, and an expectation error ξ_t^{ks}. That is, I assumed that realized values of future payoffs are noisy measures of workers' expectations, following Scott (2013) and Kalouptsidi et al. (2020).[76] Note that I do not necessarily require workers to have rational expectations with regard to future value functions. Instead, the expectation error can reflect biased beliefs but needs to be mean-independent of the instruments for the components of flow utility discussed later. An alternative method of ensuring this mean independence is demonstrated by Diamond et al. (2017) in an urban location choice setting: they exploit a quasi-random assignment of renters into control and treatment groups to difference out any bias in future expectations. In contrast, I will rely on the exclusion restrictions of my instruments to identify changes in city characteristics that are likely to be uncorrelated with the expectation error, as I discuss in the identification section below.

Applying renewal actions. We can now use the renewal action assumption to difference out the future value functions – following the intuition discussed in the example shown in Figure 12. In the aggregate, this corresponds to comparing the *probabilities* that workers choose different migration paths that end up in the same location. Relating these differences in choice probabilities to differences in the characteristics of the cities lived in along the way, allows us to identify the location choice parameters.

Assume that in period $t+1$ two workers end up in the *same* location $k = l$, which one worker already moved to in period t, and that the other worker spent period t staying in city i.[77] Substituting for the ex ante value functions from Equation (15), the difference in choice probabilities then simplifies to

$$\theta_s \ln\left(\frac{\mu_{st}^{ik}}{\mu_{st}^{ii}}\right) =(U_{kst} - U_{lst}) - (\tau_s^{ik} - \tau_s^{ii}) - \beta(\tau_s^{kk} - \tau_s^{ik}) - \theta_s\beta\ln\left(\frac{\mu_{s,t+1}^{kk}}{\mu_{s,t+1}^{ik}}\right) + \xi_t^{kis} \qquad (23)$$

where $\xi_t^{kis} = \xi_t^{ks} - \xi_t^{is}$. This expression no longer depends on the future value terms which have been differenced out by the renewal state assumption. Kalouptsidi et al. (2020) note that this expression can be viewed intuitively as an intertemporal Euler equation for the optimal choice of moving probabilities. There is a clear analogy to an indifference condition, where a representative agent is trading off the probability of reaching the renewal state along different paths.

This equation now describes the current and next period moving choices as a function of differences in flow utility, moving costs, and expectation errors. Below, I discuss how this expression can be used to identify the location choice parameters of interest.

[76] While this assumption is restrictive, it is common in applications (e.g De Groote and Verboven (2019)) because the alternative involves having to specify arbitrary assumptions about how agents expect future states to evolve.

[77] The identification does not rely on this assumption: all that matters is that the workers end up in the same location. However, assuming that one of the workers can stay in place simplifies the estimating equation.

6.3 Identification of location choice parameters

To derive estimation moments for the location choice elasticity, I need to rewrite the location choice Euler Equation (23) in terms of observable variables and model parameters. I assume that moving costs are zero for households that are not moving across cities, i.e. $\tau_s^{ii} = \tau_s^{kk} = 0$. Rearranging the location choice Euler Equation (23), and substituting the flow utility components, I obtain

$$\underbrace{\ln\left(\frac{\mu_{st}^{ik}}{\mu_{st}^{ii}}\right) + \beta \ln\left(\frac{\mu_{s,t+1}^{kk}}{\mu_{s,t+1}^{ik}}\right)}_{\mathbb{M}_{st}^{ik}\text{: Adjusted relative migration prob.}} = \theta_s^{-1} \ln\left(\frac{\mathcal{R}_{kst}^{w}}{\mathcal{R}_{ist}^{w}}\right) - \theta_s^{-1}(1-\beta)\tau_s^{ik} + \tilde{\xi}_t^{kis} \tag{24}$$

where $\tilde{\xi}_t^{kis} = \theta_s^{-1} \ln\left(\frac{A_{kst}}{A_{ist}}\right) + \theta_s^{-1}\xi_t^{kis}$ and $\mathcal{R}_{kst}^{w}$ is the real wage. Here, I have gathered the unobservable amenity differences and expectation errors into the error term.

Note that we can use this equation to illustrate how the dynamic approach compares to a static location choice model. For example, if an economic recovery is expected to raise future wages in some city i, but this increase takes place over multiple years, this might bias static estimates of location choice parameters. A static model might infer that location choice is very sensitive to wage changes if many workers move into i during the period when the initial wage changes are still small. However, the dynamic model takes into account that the future value of city i has also gone up as future wage increases are anticipated. This is reflected in the expression above by accounting for increased migration flows into city i in the future, after wages have risen. The dependent variable in the dynamic model would reflect that both increases in current wages and in the future value of being in city i can be causing current migration flows – and that the true sensitivity of migration to observable wages is therefore not as large as the static parameter estimates would suggest.

Addressing unobservable amenities. An important identification concern is that the difference in amenities $\ln\left(\frac{A_{kst}}{A_{ist}}\right)$ in the error term likely affects the long-run differences in the levels of wages and house prices. This would be the case in the model presented above, as well as in most standard Rosen-Roback style spatial equilibrium models.[78] As I do not observe amenity differences, this would cause omitted variable bias in a simple OLS regression estimation of Equation 24.

I address this issue in three ways: First, I eliminate the constant components of amenities – and the constant moving costs – by first-differencing the estimating equation. This means that identification of the coefficients on wages and house prices only requires that *changes* in amenities year-to-year are not correlated with contemporaneous year-to-year changes in real wages. While there is a substantial literature finding endogenous responses in city amenities to demographic changes,[79] many amenities, such as natural advantages, climate, and cultural factors are likely to change relatively slowly – and therefore will be differenced out by this transformation.

Second, I model endogenous amenity changes as a function of local observable characteristics. I follow an approach similar to that in Diamond (2016) and model the log of amenities as a function of the local college share of the population. Moreover, in a robustness check, I will additionally model amenities as a function of an index $Amen_{it}^{1st}$ that captures the local prevalence of consumption and leisure establishments (see Appendix Section F.4 for details.).[80] That is, in the full specification,

[78] See, e.g. Shapiro (2006); Moretti (2011); Diamond (2016)

[79] For a review, see Rosenthal and Ross (2015).

[80] I assume that establishments respond with a lag to contemporaneous demographic changes, such that these

I model the unobserved amenity change as

$$\Delta \ln A_{ist} = \beta_s^{col} \Delta ColShare_{it} + \beta_s^{est} \Delta Amen_{it}^{1st} + \epsilon_{ist}^{\Delta A}., \tag{25}$$

where $\beta_s^{est} = 0$ in the main specification. This means that the baseline estimating equation becomes

$$\Delta \mathbb{M}_{st}^{ik} = \theta_s^{-1} \Delta \ln \left(\frac{\mathcal{R}_{kst}^{w}}{\mathcal{R}_{ist}^{w}} \right) + \theta_s^{-1} \beta_s (\Delta ColShare_{kt} - \Delta ColShare_{it}) + \Delta \tilde{\xi}_t^{kis,Col}. \tag{26}$$

where $\Delta \tilde{\xi}_t^{kis,Col} = \theta_s^{-1} \left(\epsilon_{kst}^{\Delta A} - \epsilon_{ist}^{\Delta A} \right) + \theta_s^{-1} \Delta \xi_t^{kis}$. Now, the residual $\Delta \tilde{\xi}_t^{kis,Col}$ only contains changes in the difference in amenities between cities that are not captured by the change in college shares, as well as changes in expectation errors.

Third, I use instruments to identify variation in wages and housing costs that is plausibly exogenous with regard to the changes in amenities and expectation errors contained in the error term. While canonical applications of this dynamic mobility estimation approach in the trade literature to occupational switching (Artuç et al., 2010) and migration (Caliendo et al., 2019) have used past values of migration shares and wages to instrument for their future values, this would not be an appropriate identification strategy in this case. Differences in wages, migration, and house prices, both past and present, are all likely to be correlated with the difference in amenities and expectational errors contained in the error term. While taking first differences and controlling for endogenous amenity changes associated with changes in skill composition, as described above, should address some of the concern over bias from unobservable amenities changes, it is still possible that "speculative" changes in expectations for local house prices and/or wages affect both migration and future realized values of these variables. This makes lagged values of the endogenous variables unsuitable as instruments in this setting.

Shift-share instruments. Ideally, I would be able to use an approach like the one used by Diamond et al. (2017) and Diamond et al. (2019) who exploit a quasi-natural experiment in San Francisco of differences in rent control assignment that allows them to compare treatment and control groups and thereby difference out local amenity effects. However, to my knowledge, no comparable institutional variation at the national level is available for the time period under consideration.

Instead, I identify exogenous variation in wages and house prices that is plausibly uncorrelated with idiosyncratic local expectation errors and residual amenity changes by again using an industry shift-share approach: Contemporaneous labor demand shocks driven by national trends are plausibly uncorrelated with local contemporaneous amenity changes and errors in workers' expectation of the future value function.

Here, I construct the shift-share instruments for the structural model by summing over national wage trends by industry, using year 2000 employment shares. In addition to the shift-share instruments for overall changes in city wages, the model with heterogeneous skill groups additionally suggests that we can construct skill-specific shift-shares of the form

$$B_{it,`00}^{s} = \sum_{\iota} \pi_{\iota si,2000} \Delta \ln W_{\iota,-i,t}^{US},$$

for each of the college and non-college skill groups. The only difference in construction to the shift-share instrument for the city as a whole is that the weights on the national industry trends are now

observable changes in amenities are not correlated with other contemporaneous shocks to location preferences.

given by the exposure of group s to those trends, which is proxied by the share $\pi_{\iota si,2000}$ of workers in group s who work in 2-digit industry ι in the year 2000 in city i. As before, I follow the literature in interacting these wage shocks with exogenous land constraints to identify variation in housing costs. See Appendix Section F.2 for more details on the shift-share instrument construction and identifying variation. Moreover, I construct the city-level shift shares at both a 2-digit and 3-digit industry level in order to exploit additional variation in industry wage trends within industry sectors. See Appendix Section E.4 for details on the exact instruments and moments used in this estimation – and in estimating the other structural parameters

The last issue is the measurement of the housing cost changes $\Delta \ln Q_{it}$ to be used in the estimation. As I do not have precise annual data on local rent payments that cover all U.S. commuting zones, I again rely on the mapping between changes in housing costs and house price growth through the valuation function. In the construction of real wages, I substitute $\Delta \ln Q_{it} = \Delta \ln P_{it}$ with available house price growth data and include year fixed effects in the estimation to absorb any common national changes in the valuation term α_t over time. Moreover, the fact that the location choice parameters are estimated in differences means that constant differences α_i in housing valuation across cities do not affect the estimation.

6.4 Identification of housing supply elasticities

In order to estimate the parameters of the housing supply function, I first take logs of the house price function

$$\ln P_{it} = \ln \tilde{\phi}_i + \ln \tilde{\phi}_{it} + \tilde{\phi}_i^H \ln HD_{it},$$

where $HD_{it} = (\sum_\iota \sum_s \alpha_s W_{\iota i,t} L_{is,\iota,t})$ is total housing expenditure. Then, I parameterize the elasticity with regard to changes in housing demand as a function of the unavailable share of land in the city:

$$\tilde{\phi}_i^H = \psi_H + \psi_H^l x_i^{land}$$

Moreover, I assume that the time-varying local housing productivity and housing valuation term $\ln \tilde{\phi}_{it}$ can be modeled as

$$\ln \tilde{\phi}_{it} = \eta^r \ln r_t^{mtg} + \eta^e \ln ei_t + \eta_i^H \cdot t + \epsilon_{it}^P,$$

where r_t^{mtg} are nominal mortgage rates, ei_t is the expected inflation rate, $\eta_i^H \cdot t$ is a city-specific linear housing valuation trend, and ϵ_{it}^P is a local house price residual.

I substitute this expression into the housing supply equation and take first differences to eliminate the constant component $\ln \tilde{\phi}_i$ of house price differences between cities. The estimating equation therefore becomes

$$\Delta \ln P_{it} = \eta^r \Delta \ln r_t^{mtg} + \eta^e \Delta \ln ei_t + \eta_i^H + (\psi_H + \psi_H^l x_i^{land}) \Delta \ln HD_{it} + \Delta \epsilon_{it}^P. \quad (27)$$

Again, we should be concerned that $\Delta \epsilon_{it}^P$ is correlated with housing demand. For instance, in periods when house prices are unusually high, housing demand and local population are likely to be lower as fewer workers choose to move to the more expensive city. Therefore, we need instruments for housing demand that are uncorrelated with idiosyncratic local house price shocks.

Housing demand instruments. Based on the wage and population components of housing demand, we can again use Bartik wage shift-shares as instruments to identify exogenous variation in housing demand. The identifying assumption to obtain consistent estimates of ψ_H and ψ_H^l in this case is that unobservable shocks to house prices are not correlated with economic shocks based on national trends to the city, other than through the effect of the latter on housing demand and in-migration.[81]

6.5 Identification of labor demand parameters

In order to identify how wages respond to changes in employment size – the degree to which diminishing marginal returns are balanced by agglomeration effects – I need to estimate η^{LD} in the labor demand equation

$$\Delta \ln W_{i\iota t} = \eta^{LD} \Delta \ln L_{i\iota t} + \frac{1}{\sigma} \Delta \ln \tilde{D}_{i\iota t},$$

I assume that the local demand and productivity term $\frac{1}{\sigma}\Delta \ln \tilde{D}_{i\iota t}$ can be decomposed into a national industry trend component $d_{\iota t}$, a local city trend component d_{it}, and a residual local industry component $d_{\iota i t}$, in the form

$$\frac{1}{\sigma} \Delta \ln \tilde{D}_{is\iota t} = d_{\iota t} + d_{it} + d_{\iota i t}.$$

If we estimate the resulting labor demand equation, including industry×year and CZ×year fixed effects, given by

$$\Delta \ln W_{i\iota t} = \eta^{LD} \Delta \ln L_{i\iota t} + d_{\iota t} + d_i + d_{\iota i t},$$

we can recover the labor demand parameter η^{LD} under the assumption $Cov(\Delta \ln L_{i\iota t} d_{\iota i t}) = 0$ using fixed effects OLS. Of course, it is possible that this assumption does not hold exactly. Lacking an adequate instrument for local industry employment, I will therefore pay close attention to how the estimate changes as control variables are added to see if unobservable variation might play a large role in determining the coefficient estimate.

6.6 Identification of industry choice elasticity

The elasticity of workers' industry choice $\pi_{is\iota t}$ with regard to changes in wage differentials between industries is captured by the elasticity parameter a. I can derive an estimating equation for this parameter by taking log differences of Equation 9 to obtain

$$\Delta \ln \pi_{is\iota t} = \alpha_{ist} + \alpha_{s\iota i} + a \Delta \ln W_{i\iota t} + u_{is\iota t}, \tag{28}$$

where $\alpha_{ist} = \Delta \ln \left(\sum_{\iota=1}^{N_{ind}} W_{i\iota t}^a \epsilon_{s\iota i} \right)$, and a $u_{is\iota t}$ is a stochastic error term that captures any mismeasurement of wages. This means I can estimate a as the coefficient from a regression of changes in the local industry employment shares for each education group s on local log industry wage changes, controlling for a group-specific city trend α_{ist}.

Moreover, in the empirical implementation I also condition on CZ-by-industry fixed effects $\alpha_{s\iota i}$ to control for long-run local industry trends that might be invalidating the exclusion restriction, e.g. the technology industry in San Francisco experiencing faster employment growth over the sample period in general in a way that is unrelated to wage incentives.

[81] See Appendix Section E.4 for details on the exact instruments and moments used in this estimation.

From the model, national changes in industry demand represent shocks to labor demand that are plausibly exogenous with regard to local labor supply decisions. Therefore, I can use the national leave-one-out wage growth $\Delta \ln W^{US}_{\iota,-i,t}$. in industry ι as an instrument for local industry wage growth being higher relative to the city i average. The elasticity a is therefore only identified off within-city differences in a given year in the degree to which local industries outperform their usual local industry wages as a result of national trends.[82]

6.7 Additional Data for the Structural Estimation

In addition to the data used in the reduced-form estimation that are described in Section 3.6, the structural parameter estimations require additional data inputs and transformations that are described below.

Combining low-population geographic units. As the pairwise migration data by education group that are used in the location choice regressions suffer from very small sample sizes if there are few data points per commuting zone, I improve the reliability of the estimates by combining 1990 commuting zones that have less than 50,000 residents with adjacent CZs until the combined area contains at least 50,000 residents (see Appendix Section F.5 for details). For consistency, all the structural parameters are estimated using these "Adjusted CZ" geographic units, and these are also the units used in simulating the counterfactuals.

Migration probabilities by education group. The quantitative estimation allows for different worker types to have skill group-specific moving probabilities and preferences. I proxy for skill using education and distinguish a "high" education group of workers with at least a 4-year college degree or equivalent, and a "low" education group of non-college workers.

Given enough data, we could simply calculate the empirical choice probabilities for each state of interest where the state space consists of worker types, origins, destinations, and years. Unfortunately, the IRS data used in the reduced form estimation does not allow for distinguishing flows of different groups of workers. The best publicly available migration data at an annual frequency for the U.S. that also contains migrant characteristics, is from the American Community Survey (ACS) for 2005-2017, with sample sizes only in the single-digit millions for most years.

To improve the precision of migration flow estimates by education group from this ACS microdata, I therefore apply statistical techniques for data smoothing and imputation that allow me to make use of additional information contained in the ACS data, combine information across units and years, and incorporate information from the IRS migration data. For example, I use the fact that migration by education needs to fulfill adding-up constraints and is persistent over time to impute the education share of flows using a LASSO estimator on information from other parts of the ACS data. Then, I apply an empirical Bayes shrinkage estimator (Morris, 1983; DuMouchel and Harris, 1983; Gelman et al., 2013) to combine these predicted education shares with the raw observed values. Intuitively, this adjusts the raw non-college shares by moving them towards their expected value - "shrinking" the deviation - and does so to a greater degree if the raw estimate was based on a smaller sample size. Details on the full methodology can be found in Appendix E.1.

As a result of this smoothing approach, I am able to avoid location pairs dropping out of the data entirely if their flows become too small to be recorded in the sample. This captures the fact that true migration probabilities are unlikely to be precisely zero. Other recent contributions to the

[82]See Appendix Section E.4 for details on the exact instruments and moments used in this estimation.

spatial economics literature have taken similar approaches: Almagro and Domínguez-Iino (2020) also use Bayesian smoothing with data-driven priors to obtain conditional choice probabilities from noisy data.[83]

Wage growth by education. Moving choices also depend on W_{ist}, the location-specific wage for workers with and without a college education. In order to estimate group- and location-specific wage changes, I use data from the Quarterly Census of Employment and Wages (QCEW) on industry-specific wages. I combine this data with estimates of the CZ-level education share in each industry from year 2000 Census microdata (obtained from IPUMS), to predict a constant local education group-specific exposure to each industry. These exposure shares are then multiplied by the average local wage growth in each industry to construct the growth in education-specific wage indices in each CZ.

7 Structural estimation results

This section details the estimation results for the parameters that are estimated from the structural equations of the dynamic model. An overview of all the estimated parameters used in the counterfactual simulations is shown in Table 6.

7.1 Location choice parameter θ_s^{-1} estimates

This section discusses the structural estimates of the location choice elasticity parameter θ_s^{-1}. The baseline estimates consist of estimates of the location choice Equation 26 using OLS and IV, which are shown in Table 7. The sample for this estimation covers the years 2005-2017, as that is the period for which migration data by education group is available.

Overall, in the IV estimates I find a marginally greater sensitivity to location characteristics for non-college-educated workers with regard to real wages. The location choice parameter estimates are $\hat{\theta}_{\text{col}} = 1.61$ and $\hat{\theta}_{\text{nc}} = 2.49$. Both estimates are statistically significant, but only marginally for the college workers and highly significant for the non-college workers. To ensure that this is not an artifice of how I compute standard errors, each coefficient is reported with two different standard error estimates: the first one corresponds to clustering standard errors at the origin and at the destination CZ level. The second one estimates Driscoll and Kraay (1998) standard errors, which flexibly allow for nonparametrically estimated spatial covariance, and temporal dependence (with a bandwidth of 5 years). The statistical significance of the coefficients is the same under either of these standard error estimates.

Moreover, for both groups the IV coefficients are much larger than the OLS estimates, which suggests that there is downward bias in the magnitude of the OLS coefficients. This could arise, for instance, if an increase in amenities in a city, which is not accounted for by the college share changes, both leads workers to move to certain locations and also lowers the required money wage for doing so.

These estimates of migration elasticities with regard to real wages are comparable in magnitude to those found by Diamond (2016), as well as the ratio of of the employment and earnings effects of

[83]More generally, Dingel and Tintelnot (2020) show that modeling "granular" spatial choice data smoothly as coming from a multinomial count model can lead to better predictive properties than calibrating perfectly to the raw observations. They argue that the reason is that estimates using raw observations end up overfitting to idiosyncratic shock realizations.

a TFP shock found in Hornbeck and Moretti (2019).[84] However, these papers are not necessarily directly comparable as I am estimating a dynamic model with bilateral migration costs on annual data, whereas Diamond (2016) is estimating a static model at decadal frequencies without bilateral migration costs. The estimates that are most closely comparable are perhaps the labor supply elasticity estimates of Artuc et al. (2020), who estimate a dynamic location choice model with annual data in Brazil, and obtain a similar estimate of 1.962 for their equivalent of θ^{-1}, without distinguishing education groups.

Robustness checks for the location choice parameter estimates. In order to see how sensitive these estimates are to variations in the specification, I report a number of additional results, with IV coefficient estimates shown in Table 8 (and OLS estimates in Appendix Table A4).

The importance of trying to capture endogenous amenities by including the change in differences in college shares is explored in Columns 1 and 5 of Table 8 for college and non-college workers: when I omit the college share changes as a proxy for amenity changes, the estimated real wage elasticities are less than 25% smaller for both groups – and the change is not statistically significant. Columns 2 and 6 test an additional way of capturing variation in amenities, controlling for an amenity index,[85] but this variation does not affect the size of the real wage elasticity estimates much.

In Columns 3 and 7, I consider a specification that uses the average salary income per capita in the commuting zone, based on IRS salary data, rather than distinguishing wages by education group.[86] The effect of this change is that the wage coefficients decline in magnitude by about 60% for college workers and 35% for non-college workers, and only the non-college parameter estimate remains statistically significant. This drop in estimated effect size is intuitive as the common income measure is a noisier measure of the relevant group-specific city characteristics.

In Columns 4 and 8 of Table 8 I estimate a static version of the location choice equation that only includes contemporaneous migration choices in the dependent variable. That is, the dependent variable is $\ln\left(\frac{\mu_{st}^{ik}}{\mu_{st}^{ii}}\right)$ and is not adjusted for changes in future migration probabilities. The static IV estimate of location choice elasticity θ_s^{-1} is smaller and statistically not significant for either group. This suggests that, on average, changes in the option value of different locations are offsetting contemporaneous differences in wages.[87] These effects from changes in future values are controlled for in the dynamic model estimates in Table 7 by explicitly including changes in future migration in the dependent variable.[88]

[84]See Table 4, column 8, Panels A and B: the ratio of the employment to earnings effects is $\frac{3.35}{1.54} = 2.18$, which is squarely in between my estimates for the different education groups.

[85]See Appendix Section F.4 for details on the construction of this index.

[86]I also assume a common housing expenditure share of 65% in constructing real wages, such that the real income growth variable in this specification is the same across education groups.

[87]This would be the case, for instance, if some cities usually lead in wage changes that then propagate to other locations: when the leading city sees its wages increase, but migration flows from laggards do not increase correspondingly, this would appear in a static model as if workers are unresponsive to wage changes. However, workers in this example anticipate that the same wage increase will come to their laggard locations in the future, so the value of staying in place also increases, which lowers the probability of moving.

[88]An analogous result has been found for occupational mobility by Traiberman (2019), who shows that when future human capital accumulation is ignored, i.e. the estimation is static, the estimated elasticity of occupation choices with regard to observable wage differences is substantially smaller. The rationale given is similar to the intuition for choices between cities that I discussed above: wage differences between occupations appear larger than they really are when we do not account for the similarity in future values.

7.2 Housing supply parameter $\tilde{\phi}_i^H$ estimates

In this section, I discuss the house price elasticity parameters obtained from estimating the inverse housing supply function in Equation 27 over the 2000-2017 period.[89] The IV estimates are shown in Table 9, with the corresponding OLS results in Appendix Table A5. As shown in Columns 1 and 2 of Table 9, the average IV effect of housing demand changes on house prices is positive and significant without allowing for heterogeneity among cities, even when I control for CZ fixed effects.

In Columns 3 and 4, I allow for heterogeneity in the inverse supply elasticities by interacting the changes in log housing expenditures with the local share of land unavailable for construction. The interaction terms are positive and significant, suggesting that greater supply constraints result in a greater effect on house prices of changes in housing demand. This finding complements similar results in the literature for the effect of population changes mediated by geographic constraints (e.g. Saiz (2010); Diamond (2016)). To interpret the magnitudes, note that land constraints are scaled to lie between zero and one.

In Column 5, I additionally control for the local consumption amenity index as this might capture changes in the quality of housing. While the amenity index has a significant positive effect on house price growth, the housing expenditure coefficients are unchanged in size and significance.

The distribution of the implied inverse housing supply elasticities is shown below Table 9. The implied mean elasticity across cities of house prices with regard to housing expenditures is 1.65 in Column 5 and ranges from 1.02 to 3.12 across cities. The elasticities that are estimated for each Adjusted CZ based on the coefficients in Column 5 are the inputs used in the counterfactual analysis.

These inverse supply elasticity estimates are towards the upper end of the range of long-run inverse housing supply elasticities found in the seminal studies by Saiz (2010) and Diamond (2016). However, their analyses are not directly comparable as (1) I estimate short-run inverse housing supply elasticities from annual data where they use decadal intervals; (2) My sample includes the large increase in house price volatility during the boom-bust period of the late 2000s and early 2010s, and (3) I estimate the elasticity with regard to housing expenditure, not population.

We would expect the supply elasticity of housing to be lower (and the inverse elasticity therefore to be higher) in the short run than in the long run. As a result, estimates comparing house prices at decadal frequencies might not capture a large share of the short-run effect of population changes on house prices.

As regards the difference in sample period, greater price responsiveness in the last two decades could be the consequence of, for example, financial innovation that enabled borrowing by home buyers (Pavlov and Wachter, 2009), or an increase in credit supply (Mian and Sufi, 2009). At the same time, there is some evidence that regulatory constraints on housing supply have become more restrictive in recent years (Ganong and Shoag, 2017; Gyourko et al., 2019), which would also increase the slope of the supply curve. In fact, my inverse supply elasticity estimates are towards the *lower* end of those implied by the sluggish short-run supply responses found by Gorback and Keys (2020) for U.S. cities over the last decade.

[89] This period is shorter than the 1990-2017 time period used for the reduced-form estimation, as I only compute counterfactuals using post-2000 data and am aligning the parameter estimates with that period.

7.3 Labor demand parameter η^{LD} estimates

The OLS fixed effects estimates of the labor demand scale parameter are shown in Table 10. The first column shows the estimate without any control variables, and Columns 2 and 3 add controls for industry-year and city-year fixed effects. The estimated coefficients are remarkably stable across these specifications, representing an agglomeration elasticity of wages to local industry employment of 5.7%.

While this identification does not ensure a causal estimate, my coefficient turns out to be right in the middle of the range of benchmark estimates of agglomeration effects surveyed by Combes and Gobillon (2015), who note that typical values fall in the range between 0.04 and 0.07 (p. 299). The same survey paper also notes that correcting for endogeneity through instruments usually has little effect on the magnitude of agglomeration effect estimates.

7.4 Industry choice parameter a estimates

The sample for the estimation of Equation 28 consists of all NAICS 2-digit industries in continental U.S. Adjusted CZs over the 2000-2017 period, which is the period that most closely corresponds to the period for which I will simulate counterfactuals. In line with the timing assumptions of the mode, the industry choice shares $\pi_{is\iota t}$ for period t are computed from period $t+1$ employment shares – however, this timing assumption does not substantially affect the results.

Table 11 shows the IV estimates of a. As the education-specific employment in each industry is imputed using constant education shares by industry,[90] the changes in industry employment shares by education group are mostly driven by total employment changes – and the estimates are therefore very similar between the two education groups. Once city trends are controlled for, I find statistically significant positive effects of higher wages on employment shares. That is, workers move into industries that pay more.

The estimated wage elasticity of industry choice $\hat{a}$ in the full specification averages 0.325 across both education groups (Columns 3 and 6, weighted by their population share) – and this is the common value that I will use in the counterfactual simulations.

Comparison to the literature. There does not seem to be a consensus in the literature on the value of the elasticity of industry choices with regard to wages in the setting which I am considering, which includes no switching costs but allows for a local group-specific industry match quality.

In a UK setting, Pessoa (2016) estimates an elasticity of sector employment to the value of switching into the sector of 0.027. In contrast, in a Spanish setting with switching costs, Fuchs et al. (2018) estimates an elasticity of sector choice of 1.35. Dix-Carneiro (2014) estimates a sector choice idiosyncratic shock heterogeneity scale parameter of 2.15 in Brazil, and Ashournia (2018) finds 2.3 in Denmark, which would correspond to a wage elasticity of 0.4-0.5 in my setting. Caliendo et al. (2019) find an elasticity of switching state *and* sector in the U.S. of slightly less than 0.5 at an annual frequency. Based on this wide range of values, my wage elasticity estimate $\hat{a} = 0.325$ (corresponding to an idiosyncratic shock scale of 3.1) seems compatible with the other findings in the literature.

[90] See Appendix Section F.2 for details on the construction of employment by education group.

8 Counterfactuals

In this section, I use the parameter estimates from the previous section and simulate the quantitative model to explore the role played by mobility and supply constraints in cross-sectional U.S. house price dynamics in the aggregate. I show that increases in migration costs (i.e. declines in mobility) lead to a more dispersed distribution of house price growth across cities after an economic shock. Moreover, when mobility is lower, policies that reduce housing supply constraints have a greater impact on the distribution of house price growth in response to a shock.

This section proceeds as follows: First, I adapt the "dynamic hat algebra" approach of Caliendo et al. (2019) to my setting with heterogeneous workers, to be able to simulate counterfactual time series.[91] I show that the key equations of my model can be rewritten in changes over time, which implies that knowing the *changes* in amenity and productivity fundamentals is sufficient to solve for population and house price changes starting from observable values in a baseline period – that is, I require no knowledge of the *levels* of unobserved fundamentals.

Next, I address the issue that, in models with migration costs, due to the adjustment frictions, observed outcomes do not necessarily represent steady states. To compare paths of the economy under different counterfactual scenarios, I therefore first need to infer the path to a steady state equilibrium that would be consistent with the observables under the baseline scenario. I adapt the solution method of Ahlfeldt et al. (2020) to show that it is possible to infer stationary steady states consistent with the distribution of fundamentals at a particular point in time by simulating the path of this economy under the assumption of a convergent path of unobservable fundamentals until the population allocations asymptote.

After that, I discuss how I resolve a practical limitation of U.S. migration data: Migration flows by education group are only available from the American Community Survey for the years after 2005. However, to study the housing boom of the early 2000s, I need to set the starting point for counterfactual simulations at an earlier year. I show how to overcome this data limitation by using observable outcomes and migration aggregates to infer education-group-specific migration flows in the years before 2005 that are consistent with the aggregate outcomes observed. This method can yield imputed migration flow matrices by education group for the pre-2005 years, which I then use as the input for counterfactual scenarios starting in the year 2000.

Last, I use this counterfactual steady state equilibrium method to compare simulated transition paths for house price growth in response to industry wage shocks and changes in fundamentals. In particular, I explore the role that migration costs, housing supply constraints, and their interaction, play in compressing or widening the distribution of house price growth across cities.

8.1 Rewriting the dynamic model in changes

This section shows how to rewrite the dynamic model in changes, such that it can then be simulated based on series of growth rates without having to know the levels of unobservable fundamentals. This derivation builds on the method used in Caliendo et al. (2019), which I adapt to the case of heterogeneous migrant groups. The full details of the derivation can be found in Appendix Section G.1.

First, by taking the ratio of migration probabilities from equation 14 between two time periods

[91] The approach of Caliendo et al. (2019) builds on Dekle et al. (2008) who show in a static model that the observed allocations are sufficient statistics for unobservable fundamentals.

I can write

$$\mu_{s,t+1}^{ik} = \frac{\mu_{st}^{ik}(\dot{\mathcal{U}}_{t+1}^{ks})^{\frac{1}{\theta_s}}}{\sum_j^N \mu_{st}^{ij}(\dot{\mathcal{U}}_{t+1}^{js})^{\frac{1}{\theta_s}}} \tag{29}$$

which uses the notation $\mathcal{U}_t^{is} = \exp(v_t^{is})$, and $\dot{x}_{t+1} = \frac{x_{t+1}}{x_t}$.

Similarly, by differencing the conditional value function from Equation 12 and substituting for the option value and flow utility, I can write the change in the value function as

$$\dot{\mathcal{U}}_{t+1}^{is} = \dot{A}_{is,t+1}\dot{\widetilde{W}}_{is,t+1}(\dot{Q}_{i,t+1})^{-\alpha_s}\left(\sum_k^N \mu_{s,t+1}^{ik}(\dot{\mathcal{U}}_{t+2}^{ks})^{\frac{1}{\theta_s}}\right)^{\beta\theta_s} \tag{30}$$

Here, I have also assumed rational expectations on the part of the migrants to drop the expectation operators for future value terms. Rewriting the changes in wages, house prices, industry employment, and local populations in similar ways, I obtain a system of equations that can be solved forward, starting from a particular period's values. The assumption necessary to do so is that the changes in unobserved fundamentals along the simulation path are known. That is, I assume that the capitalization of rents into house prices is stable and that we are given a convergent series of changes in amenities and productivity which eventually asymptotes to zero growth in these values in the long run steady state. Under these conditions, I solve for steady state transition paths under different assumptions about fundamentals.

8.2 Stationary spatial equilibrium with forward-looking location choices

This section builds on Ahlfeldt et al. (2020) and shows how to infer stationary spatial equilibria in forward-looking spatial equilibrium models.

Equilibrium definition. I define a stationary spatial equilibrium in this setting as occurring at time T, if for all $t > T$ population sizes for each education group in each location are stable over time. That is, it requires

$$\dot{L}_{is,t+1} = 1 \;\forall\; i, s, t > T.$$

This is equivalent to requiring that outflows and inflows of each population group are balanced (Ahlfeldt et al., 2020). Note that this does not require that there are no migration flows. Due to idiosyncratic location amenity shocks, there may still be gross migration between locations.

I define the "steady-state stationary equilibrium" (SSE) associated with a particular time period t's observed allocation of workers and prices to be the stationary equilibrium that is attained, if, starting at t, all fundamentals (i.e. amenities, demand shocks, house price valuation function changes) were held fixed at their period t values forever, while population, wages, and prices converged to their long-run levels.

While this does not constitute formal proof, congestion forces are large based on my parameter estimates above, with house price responses to housing demand being much stronger than wage increases as a result of agglomeration. This makes it likely that there is a unique stationary equilibrium that corresponds to each observed non-stationary allocation. The proof of this assertion is left for a future iteration of this project.

Solving for the steady-state stationary equilibrium To find the stationary equilibrium consistent with a particular time period's observables, I use the dynamic hat equations derived in

Section 8.1. In particular, the SSE for period t corresponds to the outcomes of a simulation of the model in changes, where the fundamental paths have all been set to immediate convergence. That is, it corresponds to solving the equations forward under the assumption that

$$\{\{\dot{A}_{is,t'+1}, \dot{\tilde{D}}_{i\iota,t'+1}\}_{s=1}^{\mathcal{S}}\}_{i=1}^{N} = 1 \ \forall t' > t.$$

In addition, the dynamic forward-looking model requires us to impose an assumption that the future utility ratio $\dot{\mathcal{U}}_{t+2}^{ks}$ converges to one, i.e. no growth, at some future period.[92]

Under these assumptions, the economy settles into its SSE after a couple of years, starting from its period t allocation. The exact algorithm that I use to solve for this SSE is detailed in Appendix G.2. Given the focus of this paper, I will focus on differences in city-level house price growth along different equilibrium paths as the main outcome of interest.

An important requirement to be able to compute the steady-state equilibria for a year's economy is that we have data on observables and migration flows by education group for the reference year. In Section 8.3, I discuss how to impute the migration flows for reference years in past time periods from time series of observables and aggregate migration flows.

8.3 Imputing migration flows by education group for historical time periods

U.S. migration flows by education group are only available for the limited range of years 2005-2017 from the ACS. However, to compare the effect of different parameters on steady states around the period of the housing boom of the early 2000s, I would like to set the starting point for counterfactual simulations in the year 2000. I overcome this data limitation by using observable aggregate outcomes (population, wages, and house prices), together with known migration aggregates from the IRS, to solve for model-consistent education-group specific migration flows in 2000-2004. The obtained migration flow matrices by education group for earlier years can then be set as the starting point for comparative statics.

The exact algorithm that I use to impute past migration flows is detailed in Appendix G.3. In short, the method requires an existing time series of utility growth from some time T_1 forward, a mobility matrix at time T_1, and a full time series of within-period prices (wage and house price growth) and aggregate population allocations. I obtain the forward-looking utility growth by solving for the stationary steady state equilibrium consistent with 2005 observables, which yields an expected path of utility during the transition to equilibrium. That is, implicitly I assume that agents behave in the imputed pre-2005 periods as if the economy was going to settle into equilibrium post-2005.

I can then infer historical education group-specific migration flow matrices in 2000 by running the dynamic hat model backwards. In each backward step, I solve for the education group flows that are consistent with the utility changes implied by 2000-2004 aggregate wages, population and house price growth, and gross aggregate outmigration in each city. the resulting flow probabilities for the year 2000 by education group are then treated as data in simulating counterfactuals that start in that year.

[92]This assumption rules out scenarios akin to "rational bubbles" in location values, where accelerating future growth expectations can justify continued movement towards particular cities without reaching a steady state even in the face of strong congestion forces. In practice, I set convergence to occur 100 periods after the time period of observation for which the steady state is computed - however, in practice the steady state values stabilize much earlier than that.

8.4 Model fit: overidentifying restrictions

The imputation of historical migration flows by education group in the year 2000 from 2005 data is a good test of the model's fit to the data. Therefore, in this section I use year 2000 moments that were not targeted in the imputation as overidentifying restrictions for assessing the model.

Imputed college shares vs. actuals. To see how closely the flows by education group map to actual outcomes, I first compute the college share by CZ in 2000 which is implied by reversing the imputed migration flows for each education group for 2000-2005, starting with 2005 ACS populations by education by CZ. Then, I compare this imputed college share to an estimate of the *actual* college shares by CZ from the Year 2000 Census 5% IPUMS sample. Note that there are several reasons why these two college share estimates might not be the same – even if I had access to the true migration flow matrix by education: On the one hand, there may be sampling bias and noise, as I am comparing an estimate from a Census sample (not the population) to an ACS sample estimate, both of which are likely to be measured with some noise due to the limited sample sizes. On the other hand, my model does not account for the extensive margin of the change in college-educated population. That is, over this time period the share of college-educated workers in the U.S. population is increasing, whereas my model assumes that a fixed college-educated population moves between cities.[93] Also note that I did not impose any information on population by education group or bilateral flows on the imputation procedure during the 2000-2004 years to which it is applied.

Nonetheless, the imputed model flows do a respectable job of predicting college shares by CZ: The left panel of Figure 13 plots actual Census 2000 college shares by Adjusted CZ over college shares implied by the model-based migration flow imputation. The model fits the data quite closely: the relationship shown in the graph has a significant positive slope and a linear regression R^2 of 37%.

Imputed relative migration flows vs. actuals. More important for the geographic differences in outcomes is whether the model imputes the relative size of migration flows correctly. To assess model fit in this dimension, I compare the total city-to-city migration flows, obtained by aggregating the education-specific imputed flows, to the observed totals in the aggregate IRS migration data for the year 2000 – an implied outcome that was not a target in the imputation. The right panel in Figure 13 compares the log of total outmigration shares in the IRS data to the value obtained by summing across the model-imputed flow shares applied to imputed population totals by education group for the year 2000. As the graph shows, there is a strong positive relationship, with an R^2 of 0.77 and a linear fit slope of 0.8.

Another measure of the ability of the model to predict relative flows is the correlation between the imputed outflow share or the relative rank of each destination city from the perspective of an origin city and the actual data on these relative importance measures. I compute these correlations for each origin CZ and report the distribution of the CZ-level correlations in Table 12. The median CZ's correlation between actuals and the model imputed destination flow shares and ranking are 80% and 95%, respectively. In general, the model imputed flows have almost universally high correlations with origin CZs' actual flows to other cities, which is reassuring with regard to the model's ability to match the relative size of spillovers.

[93] I do account for general population growth by allocating IRS population counts without a migration origin to cities, but in doing so I assume that these exogenous flows leave the college share of their destination cities unchanged.

8.5 Baseline wage shock construction

In this section, I construct national industry wage shocks for the boom periods 2000-2007 and 2012-2017 and evaluate their baseline effect on the distribution of house price changes across cities. The effect of these shocks will later serve as my benchmark for evaluating how house price dynamics change in the counterfactual scenarios. Throughout, I will be comparing model-predicted transition paths towards stationary steady state equilibria.[94]

Wage shock construction. I estimate industry wage trends for different time periods for each 2-digit NAICS industry and year, controlling for city-level differences in wages, by estimating the following regression (weighting by employment) for 1990-2017:[95]

$$\ln W_{i,\iota,t} = \theta_i + \theta_{\iota t} + \epsilon^w_{i\iota t}$$

To compute the "industry wage shocks" for the boom periods 2000-2007 and 2012-2017, I then take the average of the estimated industry log wage changes $\Delta\theta_{\iota t}$ over each period. This average industry growth is then applied to each industry-by-city unit in the model for a number of years that corresponds to the period in question (e.g. applying the average 2000-2007 inudstry wage growth for 7 years starting in 2000 in the model).

To visually assess the spatial disparities in the initial impact of this wage shock, I compute the approximate overall effects at the city level by weighting the industry shocks by the local industry employment shares in the baseline years 2000 and 2012 (similar to a shift-share wage shock). The predicted total nominal wage growth for each city over the two periods is illustrated in Figure 14. The maps in Panels A and B show the relative size of the shocks for all Adjusted CZs in both boom periods.

Importance of wage shocks. As a measure of how important these wage shocks are relative to other city-level shocks occurring during the housing boom periods, I first compare the predicted effect of wage shocks on city-level house price growth to actual house price growth during these periods.[96] Note that in the model, and in reality, many other shocks – credit supply changes, amenity changes etc. – will be driving house price growth, so this analysis only serves to highlight the importance of this particular channel relative to others.

Panels A and B of Figure 15 show that the national industry wage trends, when filtered through the dynamic model, can explain 11-23% of the actual cross-sectional variation in house price growth during the boom periods. A 1 ppt difference between cities in predicted house prices on average corresponds to a 1.1-1.3 ppt difference in reality. Moreover, the levels of changes are predicted well by nominal wage shocks, which are predicted to generate average house price growth of 48 ppt and 20 ppt in 2000-2007 and 2012-2017, compared to 45 ppt and 16 ppt in the actual data. Overall, these numbers imply that industry wage shocks are a substantial factor in driving cross-sectional house price dynamics during these periods.

[94] The answers would be very similar when looking only at long-run steady state comparisons.

[95] Note that this measure is not constructed as leave-one-out for each city for simplicity, but given the large number of Adjusted CZs, the individual weight of each CZ in a particular 2-digit industry is very small.

[96] To be precise, I compute the difference in house price growth during the first T years of the transition towards a stationary steady state equilibrium with and without the industry wage shocks. Then, I compare that to the actual house price growth during the T year reference period for which the industry wage shocks were computed.

8.6 Counterfactual Scenario 1: Mobility changes

In this analysis, I explore the role of migration spillovers in distributing unequal economic shocks between cities by contemplating what happens to cross-sectional house price dynamics when migration costs change, making moves easier or harder.

Context: the secular decline in mobility. The effect of changes in migration costs has acquired policy relevance more recently as a result of accumulating evidence that inter-state migration rates have been declining over the last three decades (Molloy et al., 2011; Molloy and Smith, 2019). These changes in mobility affect the dynamism of the economy. For example, Dao et al. (2017) show that the responsiveness of migration to local labor market conditions has been weakening since the early 1990s. They show that this decline in migration elasticity is driven entirely by declines in out-migration from areas with negative shocks, consistent with workers facing a greater cost of leaving their current places of residence. The causes of these trends are not well established: Molloy et al. (2017) show that changes in the demographic characteristics of the U.S. population leave a significant share of the trend in mobility unexplained. They also argue that the decline in geographic mobility is in part driven by a decline in the rate at which workers change jobs. One possible explanation for this reluctance to switch jobs may be that workers face an increasing risk of losing human capital as a result of career changes if they become unemployed (Fujita, 2018).

Another explanation for declining geographic mobility is that the families of historical migrants who moved to fast-growing cities in the American West and South have become more rooted as successive generations grow up and call those areas their home (Coate and Mangum, 2019). Thus, as a larger share of the U.S. population lives close to their birthplace and their families, those local ties increase the cost of moving away (Zabek, 2019). In addition, young adults have been increasingly likely to live with their parents, which reduces their likelihood of moving and can explain a substantial share of the decline in inter-state mobility for that age group (Lei and South, 2020). At the same time, migration between states has become more costly as a result of increasing regulatory and institutional barriers to transferring human capital: Johnson and Kleiner (2020) show that state-specific occupational licenses reduce inter-state mobility of licensed workers by 36%. Furthermore, increasing differences in health care markets between states may have increased the risk of losing benefits when moving across state lines (Schleicher, 2017). To the degree that migration spillovers are an important adjustment mechanism in the U.S. economy, this decline in mobility could affect the dynamics of the economy, including housing markets, which is what I quantify in this section.

Mobility scenarios. For the purposes of this counterfactual, I will consider the effect of exogenous changes in structural migration costs. To be precise, I will assume that the bilateral migration cost component $\tilde{\tau}_s^{ik} = exp(-\tau_s^{ik})^{\frac{1}{\theta_s}}$ in the moving probability expression in Equation 14 can be decomposed into an element related to the personal cost of leaving the current city, a cost of changing states, and a bilateral component containing any other bilateral moving costs (including those shown in Table 2), as follows:

$$\tilde{\tau}_s^{ik} = \tilde{\tau}_s^{ik,\text{Leave}} \cdot \tilde{\tau}_s^{ik,\text{State}} \cdot \tilde{\tau}_s^{ik,\text{Other}}.$$

Then, I consider the following scenarios: (A) The cost of leaving the current city of residence falls by a third ($\tilde{\tau}_s^{ik,\text{Leave}}$ increases by a factor of 1.5 if $i \neq k$). (B) Inter-state migration costs become prohibitively large, i.e. $\tilde{\tau}_s^{ik,\text{State}} = 0$ if i and k are in different states. (C) Migration costs to *any* other city become prohibitively large ($\tilde{\tau}_s^{ik,\text{Leave}} = 0$ if $i \neq k$). See Appendix Section G.4 for details on the implementation of these changes in the dynamic hat model.

Change in migration distances. The effect of these migration cost scenarios on the ability of shocks to spill over across distances can be visualized by plotting the distribution of average migration distances of inter-city flows. Appendix Figure A10 shows the distribution of empirical baseline migration flow distances in log miles for college (Panel (a)) and non-college workers (Panel (b)) in 2012, plotted in blue.

Scenario A (decrease in migration cost), which is not shown, would consist of shifting the entire distribution up without changing its shape, as it distributes mass from zero distance moves (not shown on the graphs) to positive distance moves. Scenario B (prohibitively high inter-state migration costs), in contrast, (plotted in red) shifts the mass of long-distance flows on the right of the blue baseline density to to the left, such that most moves become short-distance - indicated by the large spike in the red distribution in the 10-50 mile range (log distances of 2.3-4). Moreover, this change in migration costs has a greater impact on mobility for college workers because they are more likely to migrate long distances in the empirical baseline. Scenario C (no mobility) would concentrate all the probability mass at zero distance moves.

Comparing scenarios. To evaluate the cross-sectional house price dynamics under these different migration cost scenarios, I simulate the house price growth over the 2000-2007 and 2012-2017 periods on the economy's path to its stationary steady state under each assumption on migration costs. Then, I compare the simulated impact of the industry wage shocks on this house price growth trajectory in each scenario, relative to each scenario's baseline changes. This comparison to scenario-specific baselines is important because the long-run steady state of the economy changes with the migration cost fundamentals as well, and I want to distinguish the responsiveness to shocks from this change in the steady state path.

House price growth differences. How do changes in migration costs affect the impact of economic shocks on the cross-section of house prices? Intuitively, migration spillovers function as an "escape valve" that allows local shocks to spread to connected locations, transferring some of the impact from the origin to the destination city. If migration costs become sufficiently high, mobility becomes less effective as mechanism for mitigating concentrated house price impacts, because nearby cities are more likely to experience similar shocks. For instance, if displaced workers from Los Angeles, CA, are no longer able to move to Las Vegas, NV, then when Los Angeles house prices go up, migrants are limited to finding a home among nearby, and equally expensive, California cities. In general, with less mobility the impact of economic shocks will be concentrated on the housing markets in fewer cities, and differences in house price growth between cities should increase as a result.

This dynamic is shown in Figure 16, with Panels (a) and (b) showing the distribution of the effects of wage shocks on house price growth in each period under the different scenarios. As the graphs show, decreases in migration costs result in a narrow unimodal distribution of house price growth. In contrast, less mobility results in more dispersed house price growth, with a greater mass of cities experiencing lower house price growth and a fatter right tail of very high house price growth areas.

As a result, the distribution of house price effects becomes more unequal, which I confirm by computing Gini coefficients for each scenario, which are shown in Table 13. I find that in each of the time periods, lower mobility significantly increases the Gini coefficient of house price effects.

To further quantify the change in the distribution of outcomes, Table 13 shows the mean wage shock effect on house price growth under each scenario, as well as the spread in outcomes from the 75th to 25th, and from the 90th to the 10th percentile. While the average house price growth barely changes in the different mobility scenarios, the interquartile and 90th-to-10th spreads are almost twice as large in the no mobility scenario, compared to the increased mobility scenario. Comparing no mobility to the baseline scenario, the spread measures are 65-70% larger in the former scenario.

Changes in house price betas. Another perspective on the impact of mobility comes from considering what happens to the distribution of house price betas across cities as migration costs change. Panels (c) and (d) of Figure 16 show the distribution of betas for the annual house price growth effect in each period, computed by regressing each city's annual house price growth effect series (the difference of counterfactual to baseline scenario growth for each year) for the period in question on the series of leave-one-out average city effects in each year.[97] The differences between the distributions shows that a lower migration cost (higher mobility) leads to a more narrow distribution of city house price effect betas. In particular, the higher migration cost scenarios have much more mass in the far left tail, indicating that in those scenarios there is a large group of cities with house price growth that moves strongly inversely to the national average. Intuitively, higher mobility leads to more elimination of such "arbitrage" opportunities in the cross-section of cities, as workers are more likely to move to take advantage of cities becoming cheaper relative to the national average.

8.7 Counterfactual Scenario 2: Housing supply constraints

As the next variation in fundamentals, I explore the importance of differences in housing supply constraints in generating different house price growth impacts across cities. In particular, I focus on the predicted effect on house price dynamics of reducing supply constraints in the cities where housing supply is least elastic. This thought experiment is particularly salient because a number of recent papers have argued that supply constraints have played an important role in limiting economic adjustment processes in the U.S. (see, e.g. Ganong and Shoag (2017); Hsieh and Moretti (2019)).

Changes in supply constraints. I focus on two scenarios that correspond to different degrees of loosening supply constraints: reducing all house price elasticity with regard to demand changes (i.e. reducing implied supply constraints) to be no higher than (1) the 75th percentile, or (2) the median of the baseline distribution.

The effect of these scenarios on the geographic distribution of supply constraints is shown in Panels (a) and (b) of Figure 17. Each map shows in darker colors the areas where supply constraints, proxied by house price elasticities, are most reduced under the different scenarios. Under the more modest reduction to the 75th percentile (left panel), the areas where supply constraints are loosened most would be the Atlantic and Pacific coasts, as well as Florida and areas near the Appalachian Mountains. When reducing price responses all the way to the median (right panel), parts of the Far West, New England, and the Gulf Coast shore are additionally affected.

House price growth differences. These reductions in supply constraints by themselves will reduce the steady-state average level of house prices because the growing areas of the U.S. are empirically more likely to be supply constrained. However, the impact on the distribution of house price growth effects in response to shocks is ex ante ambiguous: if positive shocks were concentrated in less constrained areas, then reducing the price response in the areas that are already less affected

[97] I omit the "no mobility" scenario from the beta graphs for better visibility of the other scenarios. Zero mobility results in a narrow unimodal distribution that is almost entirely contained within the interval between beta values of -3 and 6.

by economic shocks could conceivably increase the dispersion in house price outcomes. However, to the degree that reducing the highest supply constraints necessarily reduces the dispersion of price elasticities, evenly distributed economic shocks would become less likely to have differential effects.

For the industry wage shocks, the latter effect seems to win out. The distribution of house price effects of wage shocks in the supply constraint scenarios is shown in Panels (c) and (d) of Figure 17. The simulated dispersion in house price growth outcomes is reduced significantly in both periods, with the lowest constraint scenarios (curves in green) having a narrowed distribution and a lower Gini coefficient (see Table 14) than the empirical baseline constraint scenario.

Table 14 shows additional statistics on the distribution of wage shock effects on house price growth under each constraint scenario. Comparing the baseline and the median constraint scenario, the average house price growth declines modestly by 13-15%, but the interquartile and 90th-to-10th spreads are reduced by more than 50% with lower supply constraints. This shows that reductions in supply constraints have a qualitatively similar effect on the dispersion of house price growth outcomes in response to shocks as greater inter-city mobility does.

Changes in house price betas. The impact of reductions in supply constraints on house price betas is mixed: Panels (e) and (f) of Figure 17 show the distribution of effect betas for the different constraint reduction scenarios. Generally, the shift from the baseline to median constraints seems to narrow the distribution of betas and leads to more of its mass being around a value of one, that is, indicating that cities are moving together to a larger degree. However, this shift is not monotonic as the beta distribution associated with a reduction of constraints to the 75th percentile does not fall neatly in between the baseline and median reduction scenarios. The reason why the effects of changes in supply constraints on betas is complicated is that the set of cities that is driving average national prices as a result of being very supply-constrained is changing between scenarios and so even holding many cities' dynamics constant, their betas might change between scenarios due to the differences in the series of average house price effects. Moreover, as the average effect size falls, a given house price effect will imply a larger beta - yielding a complex relationship between supply constraints and city-level co-movement.

8.8 Interaction of mobility and housing supply constraints

The two counterfactuals above showed that lower supply constraints and lower migration costs both narrow the distribution of outcomes across cities in response to wage shocks. This raises the question of how these two changes in fundamentals interact: does lower mobility increase or decrease the impact of reducing supply constraints on the divergence in outcomes? I explore this interaction by looking at the distribution of outcomes under wage shock scenarios that combine the changes in migration costs and reductions in supply constraints from Sections 8.6 and 8.7.

Combined scenario house price growth effects. The distribution of outcomes under these combined scenarios is shown in Table 15. Each of the four panels of the table looks at one statistic and shows how it changes with higher migration costs (increasing from left to right) and lower price elasticity / reduced constraints (top to bottom). Within each panel, I repeat this analysis for each time period. For instance, the lower right corner of Section B of the first panel shows a value of 17.67, which corresponds to the average change in house prices in response to 2012-2017 wage shocks in a scenario where price elasticities are reduced to the median and no migration is possible. The effect is calculated relative to the baseline path to steady state of the 2012 economy in this scenario.

In general, we can observe that the measures of the spread of outcomes (Panels 2, 3, and 4) are lowest when supply constraints are reduced the most and mobility is highest (lower left corner of each section), so the equalizing effects of both of these measures are additive to some degree. It is important to note that this is not driven by a mobility effect on the *level* of house prices: while lower constraints unambiguously lower house price levels, changes in migration costs do not have a consistent positive or negative effect at any given level of constraints. However, the effect of mobility on lowering the *dispersion* in shock effects is unambiguous.

Changing importance of supply constraints. A relevant question for policymakers may be whether potential reductions in supply constraints are more or less effective at reducing the dispersion in shock effects when migration costs are high or low. We can infer the responsiveness to housing policy by comparing the dispersion measures between the baseline and median constraint scenarios (first and third row of each section), and then seeing how this difference changes across mobility scenarios.

These effects of changing constraints on the dispersion of effects are shown in Table 16. Each entry in this table corresponds to the change from the first row (baseline constraints) to the third row (median constraints) in ppt and % in the corresponding section and panel of Table 15. For example, the first entry of -6.01 is the difference in the interquartile spread between the constraint scenarios in rows 3 and 1 of Section A, Panel 2, of Table 15. It represents how much the interquartile spread in wage shock effects is reduced by lowering supply constraints to the median in the setting where migration costs have decreased. The second entry of -51 represents the percentage change that this absolute change of -6.01 ppt represents. Going from left to right in Table 16, we can see that higher migration costs almost universally *increase* the (negative) effect that reductions in supply constraints have on the dispersion of house price growth in response to wage shocks. For example – focusing on the change in the 90th-to-10th percentile spread – the absolute decline in dispersion when supply constraints are reduced to the median is twice as large without migration.[98]

8.9 Policy implications

Declining mobility and volatile house price dynamics. As noted above, there has been a secular decline in observed U.S. inter-city mobility over the last three decades. To the degree that this declining mobility represents an exogenous change – rather than an endogenous response to house price patterns (Ganong and Shoag, 2017) – the quantitative analysis in the previous section suggests that it should result in more extreme variability in house price growth across cities in response to economic shocks. There is some evidence that this increase in variability has taken place: Panel (b) of Appendix Figure A1 plots inter-city mobility and the 90th-to-10th-percentile spread in real house price growth for each year over the last three decades. It shows that not just the average house price growth is becoming more volatile after the year 2000, but the difference between cities is also on average getting larger and more volatile. While it is beyond the scope of this study to identify to what degree the declines in mobility have been exogenous, this trend is qualitatively consistent with the effect of lower mobility on house price growth differences suggested by the model.

For policy makers, this implies that if there are policies that can lower migration costs (e.g. reductions in occupational licensing or portable welfare and unemployment benefits) then these could be used to reduce regional differences in house price outcomes and the likelihood of extreme

[98] Compare Table 16 second and fourth column, showing baseline effects of -14.11 and -5.12 in the two periods, but -28.10 and -10.72 as the effects in the no mobility scenario.

house price events.

Increasing importance of reducing supply constraints. The counterfactuals above also showed that changes in worker mobility through lower migration costs can mitigate the impact of tighter housing supply constraints that widen the gap in house price growth across cities. Conversely, an exogenous decline in mobility makes a reduction in supply constraints more effective at reducing the likelihood of extreme house price outcomes. To the degree that declines in mobility have been exogenous, this means that policy changes that reduce housing supply constraints would be *more* effective today at reducing the dispersion in house price outcomes, than in higher mobility periods in the past. That is, policy makers who want to use lower housing supply constraints to reduce the dispersion in the impact of economic shocks on cities should now more seriously consider this policy option.

Moreover, if housing supply constraints continue to tighten in the most constrained places – as has been the case over the last two decades (see, e.g. Gyourko et al. (2019); Aastveit et al. (2020)) – then enabling greater worker mobility through lower migration costs could mitigate the effect of these tighter constraints on the dispersion in house price growth across U.S. cities.

Macroeconomic policy and financial regulation. The evidence provided in this paper also has implications for macroeconomic policy and financial regulation. Glaeser (2013) reflects a common argument for why house price changes in second-tier cities were considered bubbles, when he notes that, during the run-up to the housing boom of the 2000s, "[s]ome denser, older cities like New York and Boston were doing particularly well, but that can do little to explain the boom in Las Vegas and inland California." To the contrary, the analysis in this paper shows how large coastal "superstar" cities doing well can in fact *cause* the boom in locations like Las Vegas and inland California through migration spillovers.

This interconnectedness between cities and its effect on geographic variation in house price growth is an important issue for policy makers as growth in local house prices is explicitly monitored as part of financial stability assessments in the U.S.[99] Being able to attribute local house price dynamics in part to economic shocks originating in other cities that share migration links enables a better assessment of what house price movements are related to demand fundamentals, e.g. in the form of migration flows, and which are speculative.

Affordable housing policy and urban planning. Large increases in local house prices as a result of population growth are a frequent source of acrimonious local policy debates, as rising rents displace current residents and create a backlash due to "price anxiety" (Hankinson, 2018). Attempts to accommodate newcomers through construction and increased density are often resisted by home-owners, for example by blocking rezoning attempts or setting minimum lot sizes (Glaeser and Ward, 2009). Moreover, even when permitted, construction in response to increases in housing demand usually occurs with delays, as I show in this paper (see Section 4.2.2), and as others have demonstrated.[100] As a result of such lags in construction, house prices become more volatile (Paciorek, 2013).

This volatility could be reduced and local conflicts mitigated by better predicting increases in local population. The framework in this paper for predicting spillovers from other cities' housing

[99]See, e.g. the "Financial Stability Report" by Board of Governors of the Federal Reserve System (2020), which is meant to be an "assessment of financial vulnerabilities [that] informs Federal Reserve actions" and comments on the fact that "house price-to-rent ratios vary significantly across regional markets" (p. 15).

[100]See, e.g. Oh and Yoon (2020) who highlight time delays in construction after permitting, but development lags can also arise from delays in developing marginal land (Guthrie, 2010), or from a slow permitting process (Gyourko et al., 2008).

markets based on migration links enables local housing authorities to better predict migration flows as a result of shocks in other cities. For example, migration from Boston, MA, to Portland, ME, tends to follow high house price growth in Boston.[101] Thus, if the housing authority of Portland was able to track housing markets in Boston and other cities with migration links, it could predict housing demand as a result of migration spillovers and work to accelerate construction permits and planning to get ahead of an anticipated congestion in its housing market.

More generally, my results suggest that there are externalities from local housing policies on other cities that share migration links: restrictive zoning in San Francisco is not just a local problem, but also affects Boise, ID, for instance. This highlights the need for supra-metropolitan coordination, and regulatory collaboration that takes into account the desirability of regional or national housing outcomes, not just local effects.

Real estate investors. The ability of migration links to predict house price correlations between different local housing markets is also relevant for investors exposed to residential housing markets (e.g. through residential mortgage-backed securities or real estate investment trusts). When housing markets are highly integrated, equal-weighted portfolios of housing in different U.S. regions can exhibit high portfolio risk as a result of the regions' covariance with the national cycle (Cotter et al., 2015). The mechanism proposed in this paper suggests that migration links can be a predictive tool for anticipating co-movement in housing markets, and can thereby be used to reduce real estate portfolio risk. Moreover, the model counterfactuals show how the cross-section of housing market risk can be expected to change as other fundamentals, such as mobility and supply constraints evolve.

9 Conclusion

In this paper, I have documented that house price dynamics in one city can have a causal effect on house prices in another city because of the migration connection that these cities share. Both the reduced-form and the structural estimates suggest that migration spillovers play an important role in propagating economic shocks across cities and in reducing the dispersion in house price growth effects.

Moreover, I showed that a sparsely parameterized dynamic spatial equilibrium model can generate rich patterns of geographic variation in exposure to economic shocks. The potential use of this methodology in a U.S. context is not limited to spillover effects on housing markets: the same model could be used to explore the network effects through migration channels of many varieties of shocks, such as trade shocks or natural disasters, and I aim to expand the range of economic shocks and outcomes studied in future work. For instance, it remains to be explored in more detail what we can learn about local mortgage markets and the construction sector if we consider the role of migration flows.

At the same time, the analysis in this paper has focused on the importance of *domestic* inter-city migration flows in the U.S. However, immigration to the U.S. from other countries has often been a more salient political issue than domestic flows. An important direction for future research is to extend the model in this paper to incorporate the role of international migration, which can play an important role by either complementing or counteracting the effect of domestic moves on housing markets.

[101] See Appendix Figure A11

References

AASTVEIT, K., B. ALBUQUERQUE, AND A. K. ANUNDSEN (2020): "Changing supply elasticities and regional housing booms," .

ACEMOGLU, D., U. AKCIGIT, AND W. KERR (2016): "Networks and the macroeconomy: An empirical exploration," *NBER Macroeconomics Annual*, 30, 273–335.

ADAO, R., M. KOLESÁR, AND E. MORALES (2019): "Shift-share designs: Theory and inference," *The Quarterly Journal of Economics*, 134, 1949–2010.

ADELINO, M., A. SCHOAR, AND F. SEVERINO (2012): "Credit supply and house prices: evidence from mortgage market segmentation," Tech. rep., National Bureau of Economic Research.

AHLFELDT, G. M., F. BALD, D. ROTH, AND T. SEIDEL (2020): "The stationary spatial equilibrium with migration costs," .

AHLFELDT, G. M., S. J. REDDING, D. M. STURM, AND N. WOLF (2015): "The economics of density: Evidence from the Berlin Wall," *Econometrica*, 83, 2127–2189.

ALBOUY, D. AND G. EHRLICH (2018): "Housing productivity and the social cost of land-use restrictions," *Journal of Urban Economics*, 107, 101–120.

ALLEN, T. AND D. DONALDSON (2018): "The geography of path dependence," *Unpublished manuscript*.

ALMAGRO, M. AND T. DOMÍNGUEZ-IINO (2020): "Location Sorting and Endogenous Amenities: Evidence from Amsterdam," .

ALTONJI, J. G. AND D. CARD (1989): "The effects of immigration on the labor market outcomes of natives," Tech. rep., National Bureau of Economic Research.

ANGRIST, J. D., K. GRADDY, AND G. W. IMBENS (2000): "The interpretation of instrumental variables estimators in simultaneous equations models with an application to the demand for fish," *The Review of Economic Studies*, 67, 499–527.

ARCIDIACONO, P. AND R. A. MILLER (2011): "Conditional choice probability estimation of dynamic discrete choice models with unobserved heterogeneity," *Econometrica*, 79, 1823–1867.

ARKOLAKIS, C., A. COSTINOT, AND A. RODRÍGUEZ-CLARE (2012): "New trade models, same old gains?" *American Economic Review*, 102, 94–130.

ARMINGTON, P. S. (1969): "A theory of demand for products distinguished by place of production," *Staff Papers*, 16, 159–178.

ARTUC, E., P. BASTOS, AND E. LEE (2020): "Trade, Jobs and Worker Welfare," .

ARTUÇ, E., S. CHAUDHURI, AND J. MCLAREN (2010): "Trade shocks and labor adjustment: A structural empirical approach," *American economic review*, 100, 1008–45.

ASHOURNIA, D. (2018): "Labour market effects of international trade when mobility is costly," *The Economic Journal*, 128, 3008–3038.

Autor, D., C. Goldin, and L. F. Katz (2020): "Extending the Race between Education and Technology," .

Autor, D. H. and D. Dorn (2013): "The growth of low-skill service jobs and the polarization of the US labor market," *American Economic Review*, 103, 1553–97.

Bailey, M., R. Cao, T. Kuchler, and J. Stroebel (2018a): "The economic effects of social networks: Evidence from the housing market," *Journal of Political Economy*, 126, 2224–2276.

Bailey, M., R. Cao, T. Kuchler, J. Stroebel, and A. Wong (2018b): "Social connectedness: Measurement, determinants, and effects," *Journal of Economic Perspectives*, 32, 259–80.

Bartel, A. P. (1989): "Where do the new US immigrants live?" *Journal of Labor Economics*, 7, 371–391.

Bartelme, D. (2018): "Trade Costs and Economic Geography: Evidence from the US," .

Bartik, T. J. (1991): "Boon or Boondoggle? The debate over state and local economic development policies," .

Baum-Snow, N., D. A. Hartley, and K. O. Lee (2019): "The long-run effects of neighborhood change on incumbent families," .

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

Beaudry, P., D. A. Green, and B. M. Sand (2014): "Spatial equilibrium with unemployment and wage bargaining: Theory and estimation," *Journal of Urban Economics*, 79, 2–19.

Belloni, A., V. Chernozhukov, et al. (2013): "Least squares after model selection in high-dimensional sparse models," *Bernoulli*, 19, 521–547.

Board of Governors of the Federal Reserve System (2020): "Financial Stability Report," Tech. rep.

Borusyak, K., P. Hull, and X. Jaravel (2020): "Quasi-experimental shift-share research designs," Tech. rep., National Bureau of Economic Research.

Boustan, L. P. (2010): "Was postwar suburbanization white flight? Evidence from the black migration," *The Quarterly Journal of Economics*, 125, 417–443.

Bureau, U. C. (2005–2017): "PUMS accuracy of the data," .

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.

Card, D. (2001): "Immigrant inflows, native outflows, and the local labor market impacts of higher immigration," *Journal of Labor Economics*, 19, 22–64.

Carrington, W. J., E. Detragiache, and T. Vishwanath (1996): "Migration with endogenous moving costs," *The American Economic Review*, 909–930.

Chinco, A. and C. Mayer (2016): "Misinformed speculators and mispricing in the housing market," *The Review of Financial Studies*, 29, 486–522.

Chodorow-Reich, G. and J. Wieland (2020): "Secular Labor Reallocation and Business Cycles," *Journal of Political Economy*.

Coate, P. and K. Mangum (2019): "Fast Locations and Slowing Labor Mobility," .

Combes, P.-P. and L. Gobillon (2015): "The empirics of agglomeration economies," in *Handbook of regional and urban economics*, Elsevier, vol. 5, 247–348.

Cotter, J., S. Gabriel, and R. Roll (2015): "Can housing risk be diversified? A cautionary tale from the housing boom and bust," *The Review of Financial Studies*, 28, 913–936.

Couture, V. and J. Handbury (2017): "Urban revival in America, 2000 to 2010," Tech. rep., National Bureau of Economic Research.

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.

Davidoff, T. et al. (2016): "Supply Constraints Are Not Valid Instrumental Variables for Home Prices Because They Are Correlated With Many Demand Factors," *Critical Finance Review*, 5, 177–206.

De Groote, O. and F. Verboven (2019): "Subsidies and time discounting in new technology adoption: Evidence from solar photovoltaic systems," *American Economic Review*, 109, 2137–72.

DeFusco, A., W. Ding, F. Ferreira, and J. Gyourko (2018): "The role of price spillovers in the American housing boom," *Journal of Urban Economics*, 108, 72–84.

DeFusco, A. A., W. Ding, F. V. Ferreira, and J. Gyourko (2015): "The role of contagion in the american housing boom," Tech. rep., Working Paper.

Dekle, R., J. Eaton, and S. Kortum (2008): "Global rebalancing with gravity: Measuring the burden of adjustment," *IMF Staff Papers*, 55, 511–540.

Derenoncourt, E. (2019): "Can you move to opportunity? Evidence from the Great Migration," .

Diamond, R. (2016): "The determinants and welfare implications of US workers' diverging location choices by skill: 1980–2000," *The American Economic Review*, 106, 479–524.

Diamond, R., T. McQuade, and F. Qian (2017): "The Effects of Rent Control Expansion on Tenants, Landlords, and Inequality," in *Proceedings. Annual Conference on Taxation and Minutes of the Annual Meeting of the National Tax Association*, JSTOR, vol. 110, 1–72.

——— (2019): "The effects of rent control expansion on tenants, landlords, and inequality: Evidence from San Francisco," *American Economic Review*, 109, 3365–94.

Dingel, J. and F. Tintelnot (2020): "Spatial Economics for Granular Settings," Tech. rep., National Bureau of Economic Research.

Dix-Carneiro, R. (2014): "Trade liberalization and labor market dynamics," *Econometrica*, 82, 825–885.

Donaldson, D. and R. Hornbeck (2016): "Railroads and American economic growth: A "market access" approach," *The Quarterly Journal of Economics*, 131, 799–858.

Dougherty, C. (2020): *Golden gates: Fighting for housing in America*, Penguin Press.

Driscoll, J. C. and A. C. Kraay (1998): "Consistent covariance matrix estimation with spatially dependent panel data," *Review of economics and statistics*, 80, 549–560.

DuMouchel, W. H. and J. E. Harris (1983): "Bayes methods for combining the results of cancer studies in humans and other species," *Journal of the American Statistical Association*, 78, 293–308.

Eckert, F., T. C. Fort, P. K. Schott, and N. J. Yang (2020): "Imputing Missing Values in the US Census Bureau's County Business Patterns," Tech. rep., National Bureau of Economic Research.

Eckert, F., S. Ganapati, and C. Walsh (2019): "Skilled tradable services: The transformation of US high-skill labor markets," *Available at SSRN 3439118*.

Fajgelbaum, P. D. and C. Gaubert (2020): "Optimal spatial policies, geography, and sorting," *The Quarterly Journal of Economics*, 135, 959–1036.

Favara, G. and J. Imbs (2015): "Credit supply and the price of housing," *The American Economic Review*, 105, 958–992.

Fishback, P. V., W. C. Horrace, and S. Kantor (2006): "The impact of New Deal expenditures on mobility during the Great Depression," *Explorations in Economic History*, 43, 179–222.

Fuchs, S. et al. (2018): "The Spoils of War: Trade Shocks during WWI and Spain's Regional Development," in *2018 Meeting Papers*, Society for Economic Dynamics, 1172.

Fujita, S. (2018): "Declining labor turnover and turbulence," *Journal of Monetary Economics*, 99, 1–19.

Ganong, P. and D. W. Shoag (2017): "Why Has Regional Income Convergence in the US Declined?" Tech. rep., National Bureau of Economic Research.

Gelman, A., J. B. Carlin, H. S. Stern, D. B. Dunson, A. Vehtari, and D. B. Rubin (2013): *Bayesian data analysis*, CRC press.

Glaeser, E. L. (2013): "A nation of gamblers: Real estate speculation and American history," *American Economic Review*, 103, 1–42.

Glaeser, E. L., J. D. Gottlieb, and J. Gyourko (2012): "Can cheap credit explain the housing boom?" in *Housing and the financial crisis*, University of Chicago Press, 301–359.

Glaeser, E. L., J. Gyourko, and A. Saiz (2008): "Housing supply and housing bubbles," *Journal of urban Economics*, 64, 198–217.

Glaeser, E. L., H. Kim, and M. Luca (2017): "Nowcasting the local economy: Using yelp data to measure economic activity," Tech. rep., National Bureau of Economic Research.

——— (2018): "Nowcasting gentrification: using yelp data to quantify neighborhood change," in *AEA Papers and Proceedings*, vol. 108, 77–82.

Glaeser, E. L., J. Kolko, and A. Saiz (2001): "Consumer city," *Journal of economic geography*, 1, 27–50.

Glaeser, E. L. and C. G. Nathanson (2017): "An extrapolative model of house price dynamics," *Journal of Financial Economics*, 126, 147–170.

Glaeser, E. L. and J. M. Shapiro (2001): *City growth and the 2000 census: Which places grew, and why*, Brookings Institution, Center on Urban and Metropolitan Policy Washington, DC.

Glaeser, E. L. and B. A. Ward (2009): "The causes and consequences of land use regulation: Evidence from Greater Boston," *Journal of urban Economics*, 65, 265–278.

Goldsmith-Pinkham, P., I. Sorkin, and H. Swift (2018): "Bartik instruments: What, when, why, and how," Tech. rep., National Bureau of Economic Research.

Gorback, C. S. and B. J. Keys (2020): "Global Capital and Local Assets: House Prices, Quantities, and Elasticities," Tech. rep., National Bureau of Economic Research.

Greenwood, R. and S. G. Hanson (2015): "Waves in ship prices and investment," *The Quarterly Journal of Economics*, 130, 55–109.

Gross, E. (2003): "US population migration data: Strengths and limitations," *Internal Revenue Service Statistics of Income Division, Washington, DC. http://www. irs. gov/pub/irs-soi/99gross_update. doc.*

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

Guthrie, G. (2010): "House prices, development costs, and the value of waiting," *Journal of Urban Economics*, 68, 56–71.

Gyourko, J., J. Hartley, and J. Krimmel (2019): "The Local Residential Land Use Regulatory Environment Across US Housing Markets: Evidence from a New Wharton Index," Tech. rep., National Bureau of Economic Research.

Gyourko, J., C. Mayer, and T. Sinai (2013): "Superstar cities," *American Economic Journal: Economic Policy*, 5, 167–99.

Gyourko, J., A. Saiz, and A. Summers (2008): "A new measure of the local regulatory environment for housing markets: The Wharton Residential Land Use Regulatory Index," *Urban Studies*, 45, 693–729.

Hankinson, M. (2018): "When do renters behave like homeowners? High rent, price anxiety, and NIMBYism," *American Political Science Review*, 112, 473–493.

Haughwout, A., R. W. Peach, J. Sporn, and J. Tracy (2012): "The supply side of the housing boom and bust of the 2000s," in *Housing and the financial crisis*, University of Chicago Press, 69–104.

Himmelberg, C., C. Mayer, and T. Sinai (2005): "Assessing high house prices: Bubbles, fundamentals and misperceptions," *Journal of Economic Perspectives*, 19, 67–92.

Hornbeck, R. and E. Moretti (2019): "Estimating Who Benefits from Productivity Growth: Direct and Indirect Effects of City Manufacturing TFP Growth on Wages, Rents, and Inequality," .

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

Howard, G. (2020): "The migration accelerator: Labor mobility, housing, and demand," *American Economic Journal: Macroeconomics*, 12, 147–79.

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

Jaeger, D. A., J. Ruist, and J. Stuhler (2018): "Shift-Share Instruments and Dynamic Adjustments: The Case of Immigration," Tech. rep., Working Paper.

Johnson, J. E. and M. M. Kleiner (2020): "Is Occupational Licensing a Barrier to Interstate Migration?" *American Economic Journal: Economic Policy*, 12, 347–73.

Jordà, Ò. (2005): "Estimation and inference of impulse responses by local projections," *American economic review*, 95, 161–182.

Kalouptsidi, M. (2014): "Time to build and fluctuations in bulk shipping," *American Economic Review*, 104, 564–608.

Kalouptsidi, M., P. T. Scott, and E. Souza-Rodrigues (2020): "Linear iv regression estimators for structural dynamic discrete choice models," *Journal of Econometrics*.

Kaplan, G., K. Mitman, and G. L. Violante (2020): "The Housing Boom and Bust: Model Meets Evidence," *Journal of Political Economy*, 128.

Kennan, J. and J. R. Walker (2011): "The effect of expected income on individual migration decisions," *Econometrica*, 79, 211–251.

Landvoigt, T. (2017): "Housing demand during the boom: The role of expectations and credit constraints," *The Review of Financial Studies*, 30, 1865–1902.

Lei, L. and S. J. South (2020): "The comforts of home: The association between coresidence with parents and young adults' residential mobility and migration in the United States," *Population, Space and Place*, e2323.

López-Salido, D., J. C. Stein, and E. Zakrajšek (2017): "Credit-market sentiment and the business cycle," *The Quarterly Journal of Economics*, 132, 1373–1426.

Loutskina, E. and P. E. Strahan (2015): "Financial integration, housing, and economic volatility," *Journal of Financial Economics*, 115, 25–41.

Lutz, C. and B. Sand (2019): "Highly disaggregated land unavailability," *Available at SSRN 3478900.*

Magnac, T. and D. Thesmar (2002): "Identifying dynamic discrete decision processes," *Econometrica*, 70, 801–816.

Manski, C. F. (1993): "Dynamic choice in social settings: Learning from the experiences of others," *Journal of Econometrics*, 58, 121–136.

Mian, A., K. Rao, and A. Sufi (2013): "Household balance sheets, consumption, and the economic slump," *The Quarterly Journal of Economics*, 128, 1687–1726.

Mian, A. and A. Sufi (2009): "The consequences of mortgage credit expansion: Evidence from the US mortgage default crisis," *The Quarterly Journal of Economics*, 124, 1449–1496.

Mian, A., A. Sufi, and E. Verner (2019): "How does credit supply expansion affect the real economy? the productive capacity and household demand channels," *The Journal of Finance.*

Molloy, R. and C. Smith (2019): "US Internal Migration: Recent Patterns and Outstanding Puzzles," in *Prepared for the Federal Reserve Bank of Boston's 63rd Economic Conference. "A House Divided: Geographical Disparities in Twenty-First Century America," October*, 4–5.

Molloy, R., C. L. Smith, and A. Wozniak (2011): "Internal migration in the United States," *Journal of Economic perspectives*, 25, 173–96.

——— (2017): "Job changing and the decline in long-distance migration in the United States," *Demography*, 54, 631–653.

Moretti, E. (2011): "Local labor markets," *Handbook of labor economics*, 4, 1237–1313.

——— (2013): "Real wage inequality," *American Economic Journal: Applied Economics*, 5, 65–103.

Morris, C. N. (1983): "Parametric empirical Bayes inference: theory and applications," *Journal of the American statistical Association*, 78, 47–55.

Oh, H. and C. Yoon (2020): "Time to build and the real-options channel of residential investment," *Journal of Financial Economics*, 135, 255–269.

Paciorek, A. (2013): "Supply constraints and housing market dynamics," *Journal of Urban Economics*, 77, 11–26.

Pavlov, A. D. and S. M. Wachter (2009): "Subprime lending and house price volatility," *U of Penn, Inst for Law & Econ Research Paper.*

Pessoa, J. P. (2016): "International competition and labor market adjustment," .

Piyapromdee, S. et al. (2014): "The impact of immigration on wages, internal migration and welfare," Tech. rep., Working paper.

Rosenthal, S. S. and S. L. Ross (2015): "Change and persistence in the economic status of neighborhoods and cities," in *Handbook of regional and urban economics*, Elsevier, vol. 5, 1047–1120.

Ruggles, S., S. Flood, R. Goeken, J. Grover, E. Meyer, J. Pacas, and M. Sobek (2020): "IPUMS USA: Version 10.0," .

Rust, J. et al. (1994): "Structural estimation of Markov decision processes," *Handbook of econometrics*, 4, 3081–3143.

Saiz, A. (2010): "The geographic determinants of housing supply," *The Quarterly Journal of Economics*, 125, 1253–1296.

Saks, R. E. (2008): "Job creation and housing construction: Constraints on metropolitan area employment growth," *Journal of Urban Economics*, 64, 178–195.

Schleicher, D. (2017): "Stuck: The law and economics of residential stagnation," *Yale LJ*, 127, 78.

Scott, P. T. (2013): "Dynamic Discrete Choice Estimation of Agricultural Land Use," .

Shapiro, J. M. (2006): "Smart cities: quality of life, productivity, and the growth effects of human capital," *The review of economics and statistics*, 88, 324–335.

Sinai, T. and N. Souleles (2013): "Can owning a home hedge the risk of moving?" *American Economic Journal: Economic Policy*, 5, 282–312.

Stock, J. H. and M. W. Watson (2018): "Identification and estimation of dynamic causal effects in macroeconomics using external instruments," *The Economic Journal*, 128, 917–948.

Sturm, D., S. Heblich, and S. Redding (2020): "The making of the modern metropolis: evidence from London," *Quarterly Journal of Economics*.

Sumner, S. and K. Erdmann (2020): "Housing Policy, Monetary Policy, and the Great Recession," *Mercatus Research Series*.

Traiberman, S. (2019): "Occupations and import competition: Evidence from Denmark," *American Economic Review*, 109, 4260–4301.

Zabek, M. A. (2019): "Local ties in spatial equilibrium," .

A Tables

Table 1: U.S CZ-to-CZ domestic migration links 1990-2010

Migration distance	*# of links*		*% of links that overlap after...*			
	Avg.	**Median**	**1 year**	**2 years**	**5 years**	**10 years**
Full Network	30	11	83	82	82	80
Distance > 50 **mil**	26	6	75	74	72	70
Distance > 150 **mil**	21	2	69	67	65	62

Note: Table uses continental U.S. CZ migration data constructed from IRS flows for 1990-2010 that had at least one inflow link to another CZ in 1990-2010, which includes 717 CZs. The IRS does not record flows corresponding to < 10 tax returns, and data for years past 2010 are omitted due to a methodology change in IRS gross flows post-2010. Link persistence is computed as the % of links in a given year that are still links after a certain # of years for each CZ-year, averaged over CZs and years. Networks based on migration distances (rows 2 and 3) exclude any CZ that contains counties with a centroid that is less than the stated distance away from the centroid of any county in the CZ for which inflow links are counted.

Table 2: Migration cost determinants

Dependent var.:	Log migration between cities		
Period:	1990-2017	2005-2017	
Education group:	All	College	No college
	(1)	(2)	(3)
Log distance	-1.228***	-0.682***	-0.948***
	(0.059)	(0.055)	(0.033)
Different region	-1.126***	-0.342	-0.802**
	(0.350)	(0.212)	(0.323)
Diff. region× Log dist.	0.137**	0.021	0.093
	(0.059)	(0.037)	(0.058)
Different state	-2.632***	-1.849***	-2.744***
	(0.463)	(0.406)	(0.440)
Diff. state× Log dist.	0.372***	0.187**	0.340***
	(0.096)	(0.084)	(0.086)
Nontrad. Christ. Share diff.	-0.658***	-0.389***	-0.282***
	(0.136)	(0.075)	(0.064)
Ethnicity shares vector dist.	-1.015***	-0.510***	-0.604***
	(0.170)	(0.090)	(0.098)
Water surface	-1.253***	-0.642***	-0.656***
	(0.176)	(0.131)	(0.137)
Jan. Temp. diff.	0.022***	0.006***	0.010***
	(0.004)	(0.002)	(0.003)
Industry shares vector dist.	-3.523***	-1.291***	-1.070***
	(0.302)	(0.187)	(0.181)
Observations	13,390,146	6,686,743	6,686,743
Year FE	✓	✓	✓
Origin City FE	✓	✓	✓
Destination City FE	✓	✓	✓

Heteroskedasticity-robust standard errors clustered at the origin CZ level in parentheses: * p<0.10, ** p<0.05, *** p<0.01. Analysis includes migration flows between all continental U.S. CZs, excl. New Orleans, leading to a total of 721 CZs. Migration flows by education group 2005-2017 were imputed from ACS and IRS data as described in the text. Industry share distance is computed as Euclidean distance between vectors of 2-digit NAICS employment shares in 2000. Ethnic distance is the Euclidean distance in CZ ethnicity shares in 2000. Estimated using Poisson Pseudo-Maximum Likelihood with fixed effects for origin and destination cities.

Table 3: Reduced-form IV: first-stage coefficients

Dependent variable:	Δ Network House Prices$_{it}$	
	(1)	(2)
Mig. NW Shock: Wage x Unavail. Land	0.82***	1.90***
	(0.04)	(0.18)
Mig. NW Shock: Wage	1.00***	-2.46***
	(0.02)	(0.18)
Observations	15,822	15,822
Year FE		✓
CZ FE		✓
Regional trend FEs		✓
Migration Access$_{i,t-1}$		✓
ΔW_{it}^{NW}		✓
Wage shock$_{it}$		✓
Wage shock$_{it}$ $\times$ Unavail. Land$_i$		✓

Heteroskedasticity-robust standard errors clustered at the CZ level in parentheses: * p<0.10, ** p<0.05, *** p<0.01. Includes data from 586 CZs for 1991-2017. Migration access is the migration-weighted sum of city populations. Both measures use average 1990-1995 migration flows to compute migration weights. The "Wage shocks" are Bartik shocks computed as a weighted average of national leave-one-out wage growth by industry, with weights given by local 3-digit NAICS industry shares in 1990.

Table 4: Reduced-form IV spillover estimation - baseline results

Dependent variable:	CZ Log House price growth$_{it}$					
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: OLS						
Δ Network HP$_t$	0.42***	0.26***	0.15***	0.15***	0.15***	0.14***
	(0.02)	(0.03)	(0.02)	(0.02)	(0.02)	(0.02)
Panel B: IV						
Δ Network HP$_t$	0.65***	0.53***	0.38***	0.40***	0.29***	0.23***
	(0.02)	(0.06)	(0.07)	(0.08)	(0.06)	(0.06)
Observations	15,822	15,822	15,822	15,822	15,822	15,822
1st-stage F-stat.	3,202	107	86	66	122	115
Network IV	✓	✓	✓	✓	✓	✓
Year FE		✓	✓	✓	✓	✓
CZ FE		✓	✓	✓	✓	✓
Regional trend FEs			✓	✓	✓	✓
Migration Access$_{i,t-1}$				✓	✓	✓
ΔW_{it}^{NW}					✓	✓
Wage shock$_{it}$						✓
Wage shock$_{it}$ $\times$ Unavail. Land$_i$						✓

Heteroskedasticity-robust standard errors clustered at the CZ level in parentheses: * p<0.10, ** p<0.05, *** p<0.01. Includes data from 586 CZs for 1991-2017. Network house prices are the weighted sum of other CZs' house price changes - including only CZs with > 150 mi. distance from the focal city in the migration network. Migration access is the migration-weighted sum of city populations. Both measures use average 1990-1995 migration flows to compute migration weights. The "Wage shocks" are Bartik shocks computed as a weighted average of national leave-one-out wage growth by industry, with weights given by local 3-digit NAICS industry shares in 1990.

Table 5: Reduced-form spillover effects: additional results

Dependent variable:	Δ Log House Price	Δ Log House Price	Net Migration	Log Hous. Permits	Log Mtg. Loans	Log Mtg. Lending ($)
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: OLS						
Δ Network HP$_t$	0.042	0.070***	0.019***	0.012***	0.003*	0.004**
	(0.032)	(0.023)	(0.003)	(0.002)	(0.002)	(0.002)
$\Delta HP_t \times$ Unavail. land	0.290***					
	(0.066)					
House price growth$_{t-1}$		0.395***				
		(0.027)				
Δ Network HP$_{t-1}$		0.038**				
		(0.016)				
Panel B: IV						
Δ Network HP$_t$	0.140*	-0.033	0.050***	0.007	0.023***	0.026***
	(0.073)	(0.108)	(0.016)	(0.011)	(0.008)	(0.008)
$\Delta HP_t \times$ Unavail. land	0.423***					
	(0.088)					
House price growth$_{t-1}$		0.575***				
		(0.118)				
Δ Network HP$_{t-1}$		0.156				
		(0.107)				
Observations	15,822	14,650	15,822	15,469	15,800	15,800
Year FE	✓	✓	✓	✓	✓	✓
CZ FE	✓	✓	✓	✓	✓	✓
Regional trend FEs	✓	✓	✓	✓	✓	✓
Migration Access$_{i,t-1}$	✓	✓	✓	✓	✓	✓
ΔW_{it}^{NW}	✓	✓	✓	✓	✓	✓
Wage shock$_{it}$	✓	✓	✓	✓	✓	✓
Wage shock$_{it}$ $\times$ Land$_i$		✓	✓	✓	✓	✓

Heteroskedasticity-robust standard errors clustered at the CZ level in parentheses: * $p<0.10$, ** $p<0.05$, *** $p<0.01$. Includes data from 586 CZs for 1991-2017. Migration network by distance includes all CZs for which the center of no component county is no closer than 150 miles to the center of any focal city component county. Network house prices are the weighted sum of other CZs' house price changes. Migration access is the migration-weighted sum of city populations. Both measures use average 1990-1995 migration flows to compute migration weights. The "Wage shocks" are Bartik shocks computed as a weighted average of national leave-one-out wage growth by industry, with weights given by local 3-digit NAICS industry shares in 1990. Dependent variables: Housing permits are measured in number of permitted units; mortgage loan count and lending volume includes all originated purchase loans in HMDA data; net migration is IRS domestic migration flows as a share of CZ population.

Table 6: Model parameters

Parameter	Interpretation	Source	Values
Estimated parameters			
θ_s^{-1}	Elasticity of migration choices to common city characteristics	IV estimation	1.61 (College) 2.49 (Non-college)
$\tilde{\phi}_i^H$	City-specific elasticity of house prices with regard to housing expenditure	IV estimation	1.03-3.12
η^{LD}	Agglomeration effect: elasticity of wages with regard to local industry employment	OLS fixed effects estimation	0.057
a	Elasticity of industry sector choice with regard to wage changes	IV estimation	0.325
Calibrated parameters			
β	Discount factor	Dix-Carneiro (2014)	0.95
α_s	Education-specific share of expenditure on non-tradable goods affected by local house price changes	Diamond (2016)	0.63 (College) 0.68 (Non-college)

Table 7: Location choice parameters: baseline results

Dependent variable:	Adj. Mig. Probabilities			
Education group:	College		Non-college	
Estimation:	OLS	IV	OLS	IV
	(1)	(2)	(3)	(4)
Δ Log College Real Wage diff.	0.03	1.61*		
	(0.10)	(0.82)		
	[0.06]	[0.92]		
Δ Log Non-Coll. Real Wage diff.			0.00	2.49***
			(0.07)	(0.51)
			[0.07]	[0.95]
Δ College share diff.	2.20***	-4.74	-0.59***	-6.10
	(0.31)	(6.91)	(0.18)	(4.54)
	[0.25]	[6.36]	[0.16]	[6.36]
Observations	2,877,952	2,877,952	2,877,952	2,877,952
Year FE	✓	✓	✓	✓

Heteroskedasticity-robust standard errors clustered at the origin CZ and destination CZ level in parentheses. Driscoll-Kraay standard errors in brackets: * p<0.10, ** p<0.05, *** p<0.01. Includes data from 721 CZs for 2005-2017. The instruments in all regressions consist of education-group specific wage shift-share instruments, and CZ-level shift-share instruments, as well as their interaction with CZ land constraints.

Table 8: Location choice parameters: Additional IV results

Education group:	College				No College			
Model:	No Coll. Share (1)	Amen. Index (2)	CZ Inc. per cap. (3)	Static dep. var. (4)	No Coll. Share (5)	Amen. Index (6)	CZ Inc. per cap. (7)	Static dep. var. (8)
Δ Log College Real Wage diff.	1.27*	1.61*		0.63				
	(0.72)	(0.83)		(0.53)				
Δ Log Non-Coll. Real Wage diff.					2.10***	2.54***		0.54
					(0.42)	(0.54)		(0.34)
Δ College share diff.		-5.68	-0.01	0.99		-7.21	-2.80	1.31
		(7.48)	(5.75)	(5.13)		(4.94)	(4.18)	(3.65)
Δ Amenities index diff.		-0.09				-0.08		
		(0.17)				(0.11)		
Δ Log CZ Real IRS Inc. diff.			0.65				1.63***	
			(0.59)				(0.34)	
Observations	2,877,952	2,877,952	2,877,952	3,139,584	2,877,952	2,877,952	2,877,952	3,139,584
Year FE	✓	✓	✓	✓	✓	✓	✓	✓

Heteroskedasticity-robust standard errors clustered at the origin CZ and destination CZ level in parentheses: * p<0.10, ** p<0.05, *** p<0.01. Includes data from 512 Adj. CZs for 2005-2017. The instruments in all regressions consist of education-group specific wage shift-share instruments, and CZ-level shift-share instruments, as well as their interaction with CZ land constraints.

Table 9: Housing supply: IV baseline results

Dependent variable:	CZ Log House price growth$_{it}$				
Specification:	Average Effects		Heterogeneous Slopes		
	(1)	(2)	(3)	(4)	(5)
Δ Housing Expenditure	1.64***	1.68***	1.51***	1.00***	1.02***
	(0.08)	(0.09)	(0.09)	(0.16)	(0.16)
Δ Housing Exp. × Land unavail.			0.39***	2.32***	2.32***
			(0.09)	(0.70)	(0.70)
Δ Amenities index					0.79**
					(0.35)
Observations	9,234	9,234	9,234	9,234	9,234
Pred. Inverse HS elasticities:					
Population: mean	1.64	1.68	1.61	1.64	1.65
Population: std. error			0.08	0.48	0.48
Population: min.			1.51	1.01	1.03
Population: max.			1.86	3.10	3.12
10-yr Treas. Rates & Infl. Expect.	✓	✓	✓	✓	✓
CZ FE		✓		✓	✓

Heteroskedasticity-robust standard errors clustered at the CZ level in parentheses: * p<0.10, ** p<0.05, *** p<0.01. Includes data from 512 Adjusted CZs for 2005-2017. The instruments in all regressions are city-level wage Bartik shocks and their interaction with local land unavailability and regulatory constraints. "Housing expenditures" are the log of the sum of the wage bill in each education group multiplied by the group's calibrated housing expenditure share.

Table 10: Labor demand parameter estimation

Dependent variable:	Δ Log Industry Wage$_{\iota,i,t}$		
	(1)	(2)	(3)
Δ Log Employment$_{\iota,i,t}$	0.059***	0.057***	0.057***
	(0.006)	(0.006)	(0.006)
Observations	170,802	170,800	170,800
Adj. CZ FE		✓	
Industry × Year FE		✓	✓
Adj. CZ × Year FE			✓

Heteroskedasticity-robust standard errors clustered at the Adj. CZ and industry-year level in parentheses: * p<0.10, ** p<0.05, *** p<0.01. Includes data from 512 Adjusted CZs for 2000-2017.

Table 11: Industry choice elasticity: IV estimates

Dependent variable:	Δ Log Ind. Emp. Share$_{\iota,i,s,t}$					
Education group:	College			Non-college		
	(1)	(2)	(3)	(4)	(5)	(6)
Δ Log Wage$_{\iota,i,t}$	0.059	0.297***	0.321***	0.063	0.300***	0.328***
	(0.101)	(0.112)	(0.112)	(0.102)	(0.113)	(0.112)
Observations	165,365	165,365	165,330	165,675	165,675	165,640
Year FE	✓			✓		
Industry FE		✓			✓	
Adj. CZ×year FE		✓	✓		✓	✓
Adj. CZ× Industry FE			✓			✓

Heteroskedasticity-robust standard errors clustered at the CZ level in parentheses: * p<0.10, ** p<0.05, *** p<0.01. Includes data from 512 CZs for 2000-2017. The instrument in all regressions consists of the leave-one out log change in national wages by industry.

Table 12: Correlations of total migration shares imputed by model simulation and IRS actuals for year 2000.

Measure	**Avg.**	p1	p10	p50	p90	p99
Destination rank correlation	.76	.13	.57	.80	.92	1.00
Destination outflow share correlation	.88	-.01	.71	.95	.99	1.00

Note: Table shows the correlation between the IRS migration data measure of relative destination rank or share of migration outflows and the model-imputed values. Correlations are calculated separately for each of 512 origin Adjusted CZs across all destinations with IRS migration flow information. The model-imputed flows are the value obtained by summing across the model-imputed flow shares by origin-destination pair for the year 2000, multiplied by the imputed origin population totals by education group.

Table 13: Counterfactual 1: House price effects of wage shocks in different mobility scenarios. The table below shows summary statistics for the effects of wage shocks on house prices under different mobility scenarios.

House price growth effect	Baseline mig. costs	High inter-state mig. cost	No mobility	Increased mobility
Panel A: Wage shocks 2000-2007				
Mean effect (ppt)	48.87	48.32	47.91	49.77
Diff. 75th to 25th ptl	13.26	21.00	22.02	11.74
Diff. 90th to 10th ptl	25.49	36.98	43.44	23.87
Gini coeff.	0.12	0.18	0.20	0.11
Panel B: Wage shocks 2012-2017				
Mean effect (ppt)	20.14	20.61	21.07	20.01
Diff. 75th to 25th ptl	5.83	8.04	9.06	5.22
Diff. 90th to 10th ptl	10.64	15.86	17.64	9.61
Gini coeff.	0.12	0.16	0.18	0.11

Table 14: Counterfactual 2: House price effects of wage shocks in different supply constraint scenarios. The table below shows summary statistics for the effects of wage shocks on house prices under different scenarios of reduced housing supply constraints.

House price growth effect	Baseline constraints	Reduce to P75	Reduce to P50
Panel A: Wage shocks 2000-2007			
Mean effect (ppt)	48.87	46.01	41.66
Diff. 75th to 25th ptl	13.26	11.23	6.15
Diff. 90th to 10th ptl	25.49	18.61	11.37
Gini coeff.	0.12	0.09	0.06
Panel B: Wage shocks 2012-2017			
Mean effect (ppt)	20.14	19.09	17.48
Diff. 75th to 25th ptl	5.83	4.74	2.83
Diff. 90th to 10th ptl	10.64	8.39	5.52
Gini coeff.	0.12	0.09	0.07

Table 15: House price effects of wage shocks under combinations of mobility and constraint scenarios. The table below shows summary statistics for the effects of wage shocks on house prices under different scenarios of reducing supply constraints (top to bottom) and changes in migration costs (which increase left to right). The supply constraint variations consist of no change relative to the observed baseline, as well as moving all price elasticities with regard to housing demand that are higher than the 75th percentile (median) to the 75th percentile value (median).

Mean effect (ppt)	Increased mobility	Baseline mig. costs	High inter-state mig. cost	No mobility
A: '00-'07				
Baseline supply constraints	49.77	48.87	48.32	47.91
Reduce to 75th pctl	46.91	46.01	44.68	43.57
Reduce to median	42.49	41.66	40.28	39.01
B: '12-'17				
Baseline supply constraints	20.01	20.14	20.61	21.07
Reduce to 75th pctl	19.04	19.09	19.29	19.52
Reduce to median	17.45	17.48	17.57	17.67

Diff. 75th to 25th ptl	Increased mobility	Baseline mig. costs	High inter-state mig. cost	No mobility
A: '00-'07				
Baseline supply constraints	11.74	13.26	21.00	22.02
Reduce to 75th pctl	10.26	11.23	16.87	19.27
Reduce to median	5.74	6.15	8.23	9.28
B: '12-'17				
Baseline supply constraints	5.22	5.83	8.04	9.06
Reduce to 75th pctl	4.18	4.74	6.79	7.86
Reduce to median	2.61	2.83	3.76	4.07

Diff. 90th to 10th ptl	Increased mobility	Baseline mig. costs	High inter-state mig. cost	No mobility
A: '00-'07				
Baseline supply constraints	23.87	25.49	36.98	43.44
Reduce to 75th pctl	17.36	18.61	25.45	27.74
Reduce to median	11.31	11.37	14.29	15.34
B: '12-'17				
Baseline supply constraints	9.61	10.64	15.86	17.64
Reduce to 75th pctl	7.81	8.39	10.71	11.34
Reduce to median	5.42	5.52	6.72	6.92

Gini coeff. of effects	Increased mobility	Baseline mig. costs	High inter-state mig. cost	No mobility
A: '00-'07				
Baseline supply constraints	0.11	0.12	0.18	0.20
Reduce to 75th pctl	0.08	0.09	0.13	0.14
Reduce to median	0.06	0.06	0.08	0.09
B: '12-'17				
Baseline supply constraints	0.11	0.12	0.16	0.18
Reduce to 75th pctl	0.09	0.09	0.12	0.13
Reduce to median	0.07	0.07	0.08	0.08

Table 16: Effect of supply constraint changes on the dispersion of house price effects. The table below shows changes in dispersion measures for the effects of wage shocks on house prices. The change compares the difference in the measure between the median constraint scenario and the baseline constraint scenario, under different scenarios for migration costs (increasing left to right).

Diff. 75th to 25th ptl	Increased mobility	Baseline mig. costs	High inter-state mig. cost	No mobility
A: '00-'07				
P50 - Baseline (abs.)	-6.01	-7.11	-12.77	-12.73
P50 - Baseline (%)	-51	-54	-61	-58
B: '12-'17				
P50 - Baseline (abs.)	-2.61	-3.00	-4.28	-4.99
P50 - Baseline (%)	-50	-51	-53	-55

Diff. 90th to 10th ptl	Increased mobility	Baseline mig. costs	High inter-state mig. cost	No mobility
A: '00-'07				
P50 - Baseline (abs.)	-12.56	-14.11	-22.69	-28.10
P50 - Baseline (%)	-53	-55	-61	-65
B: '12-'17				
P50 - Baseline (abs.)	-4.18	-5.12	-9.14	-10.72
P50 - Baseline (%)	-44	-48	-58	-61

Gini coeff. of effects	Increased mobility	Baseline mig. costs	High inter-state mig. cost	No mobility
A: '00-'07				
P50 - Baseline (abs.)	-0.05	-0.06	-0.10	-0.12
P50 - Baseline (%)	-45	-48	-55	-57
B: '12-'17				
P50 - Baseline (abs.)	-0.04	-0.05	-0.08	-0.10
P50 - Baseline (%)	-38	-43	-51	-54

B Figures

Figure 2: Migration inflow networks. Shading in the map indicates % of continental U.S. inflows to the city coming from the CZ over 2000-2007 period in IRS migration data.

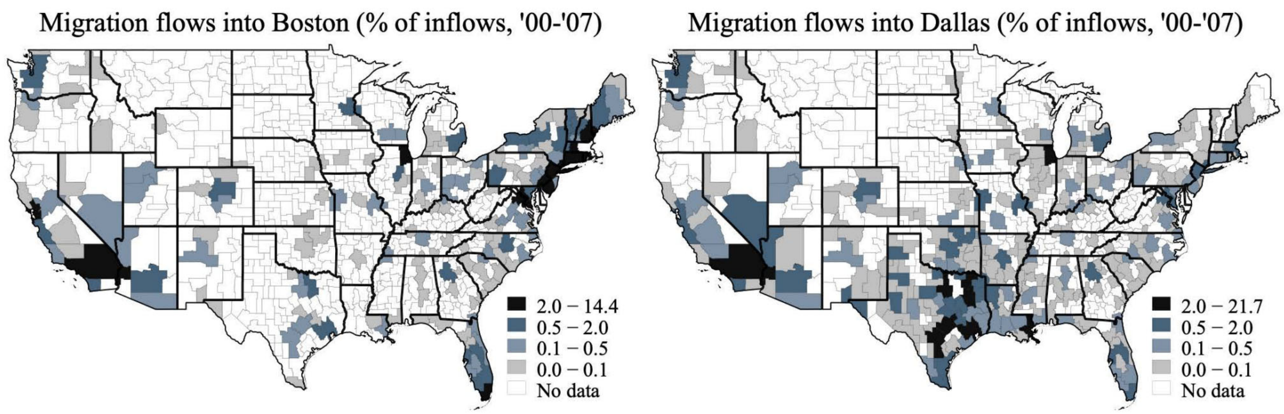

Figure 3: Population growth and house price growth. The graph plots average population growth and house price growth over 2000-2007, Data shown contains 45 CZs, consisting of all continental U.S. CZs, excl. New Orleans, with year 2000 Census population $> 0.85M$, and corresponding to $\sim 50\%$ of the continental U.S. adult population.

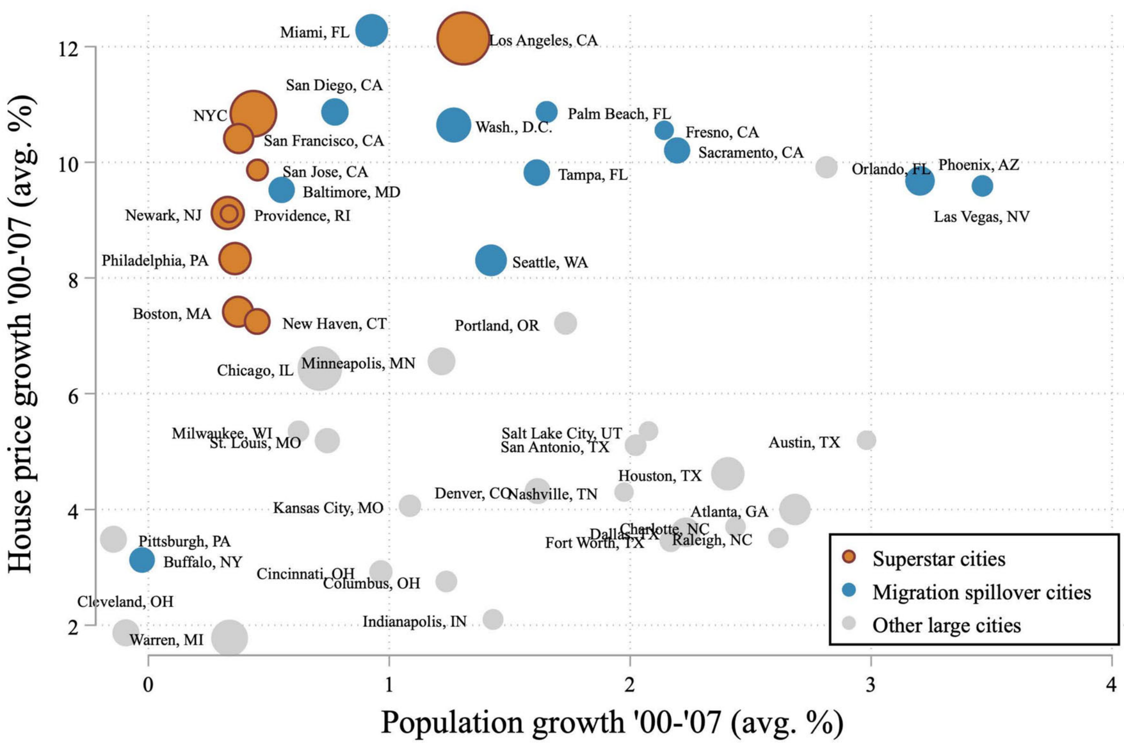

Figure 4: Ilustration of network IV identification.

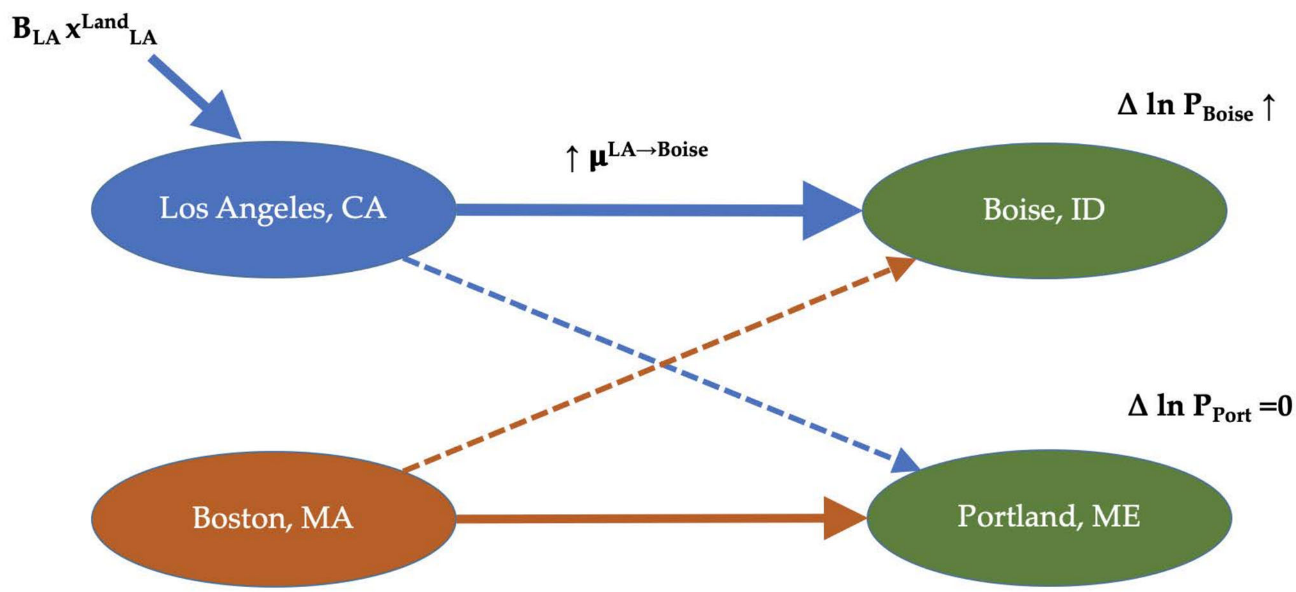

Figure 5: IV dynamic network spillover effects on house prices Graphs show IV local projection coefficients corresponding to the effect $\tilde{\eta}_h^{nw}$ of period t network house price growth on period $t-1+h$ outcome variables of the form

$$\sum_{k=1}^{h} \Delta \ln P_{i,t-1+h} = \alpha_i + \alpha_t + \tilde{\eta}_h^{nw} \Delta \mathcal{P}_{it}^{NW} + \beta'_{cum} \Gamma_{it} + \tilde{\xi}_{i,t-1+h}^{P,cum}.$$

Baseline control variables Γ_{it} in all regressions include year & CZ FEs, regional trend FEs, migration access, network log salary growth and local Bartik productivity shocks. Instruments for network house price growth in all regressions consist of the migration-weighted network averages of productivity shocks, and of migration-weighted productivity shocks interacted with other cities' land share unavailable for construction. Regressions in panels a and b also control for an interaction between local land share unavailable for construction and the productivity shocks. Regressions in panels c and d use additional instruments consisting of the baseline network instruments interacted with the unavailable land share in city i. The autocorrelation-robust regressions in panel b additionally includes one lag of local and network house price growth in addition to the contemporaneous network house price growth for which the coefficient is displayed. The lagged price variables are instrumented using two lags of the city i productivity $\times$ land share interaction, and two lags of the baseline network instruments. Estimation uses data for 1991-2017 and only inlcudes cities at $>$ 150 mi. distance in migration networks. Dashed lines show 95% CI based on std. errors clustered at the CZ level.

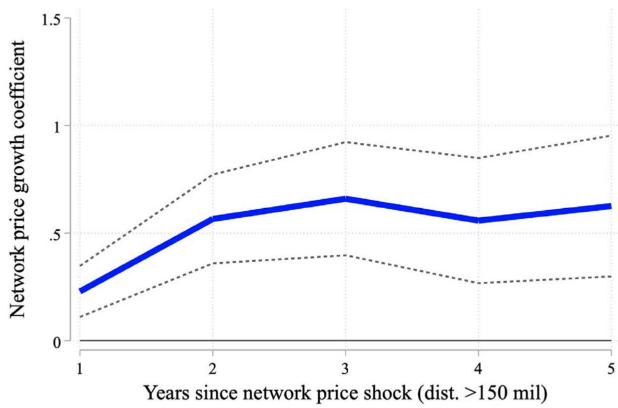

(a) Baseline effect

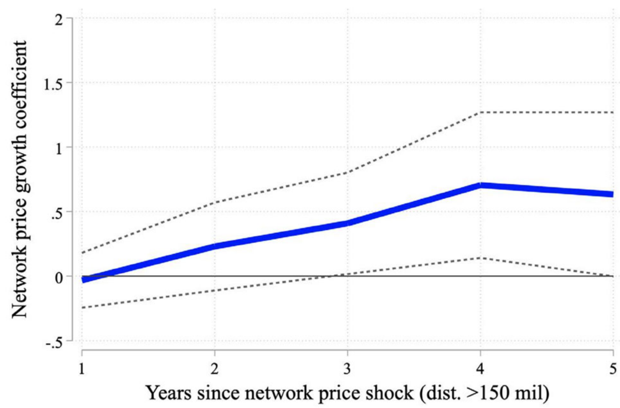

(b) Controlling for autocorrelation

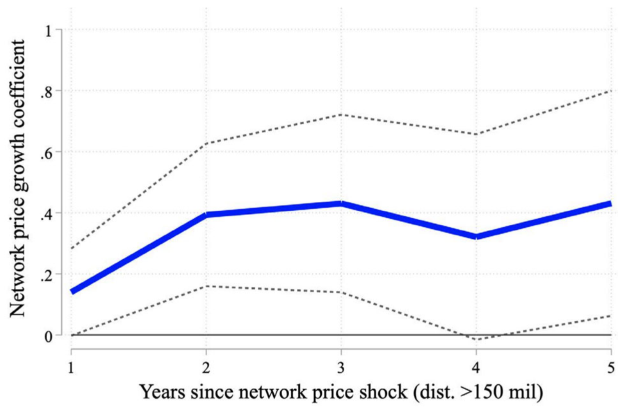

(c) Heterogeneity: uninteracted coefficient

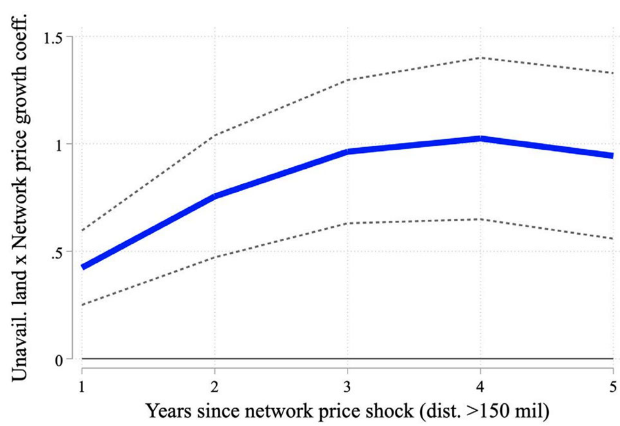

(d) Heterogeneity: interaction with land constraint

Figure 6: **Dynamic effects: other outcomes - IV** Graphs show local IV projection coefficients corresponding to the effect of period t network house price growth with period $t-1+h$ outcome variables of the form

$$Y_{i,t-1+h} = \alpha_i + \alpha_t + \tilde{\eta}_h^{nw} \Delta \mathcal{P}_{it}^{NW} + \beta' \Gamma_{it} + \tilde{\xi}_{i,t-1+h}^{P}.$$

Baseline control variables Γ_{it} in all regressions include year & CZ FEs, regional trend FEs, migration access, local log avg. salary growth, network log salary growth and local Bartik productivity shocks, as well as local productivity shocks interacted with local land share unavailable for construction. Baseline instruments for network house price growth in all regressions consist of the migration-weighted network averages of productivity shocks, and of migration-weighted productivity shocks interacted with local land share unavailable for construction. Estimation uses data for 1991-2017. Dashed lines show 95% CI based on std. errors clustered at the CZ level.

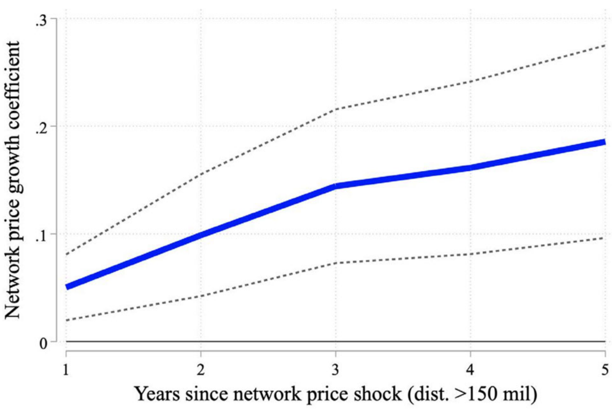

(a) Cumul. Pop Impact of Net Mig. (log pts.)

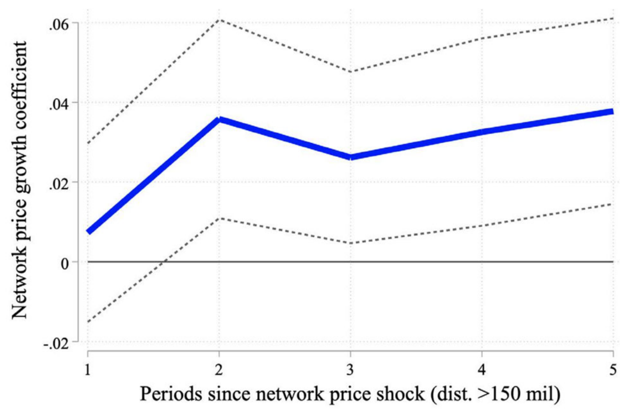

(b) Log Housing Unit Permits

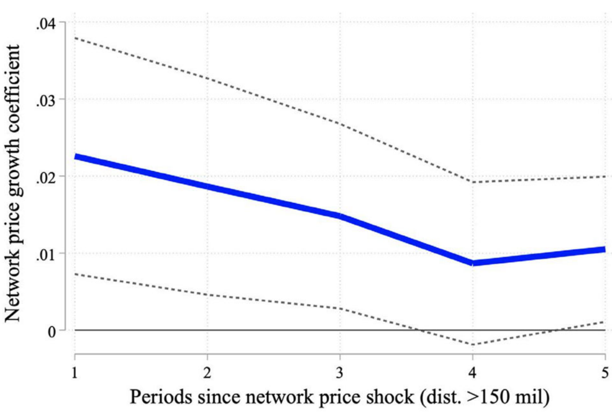

(c) Log Mortgage Purchase Loans

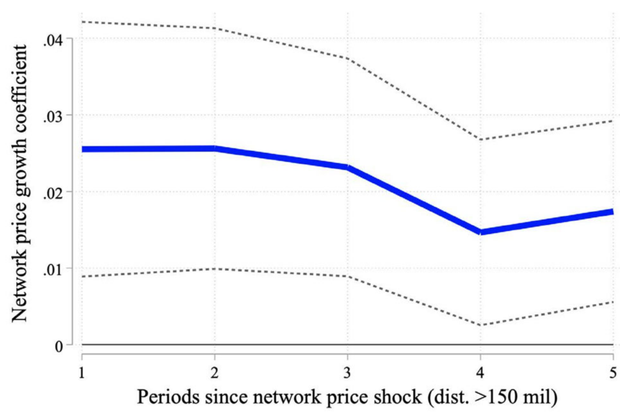

(d) Log Mortgage Lending Volume (in USD)

Figure 7: House price growth, owner-occupied vacancy rates, and time-to-build delays. The graphs in panels (a) and (b) plot CZ average annual house price growth during the housing booms of 2000-2007 and 2012-2017 over the change in average vacancy rates for owner-occupied housing. All averages are computed as differences relative to the 1991-1997 average. The graphs and fitted lines include 104 CZs for 2000-2007 and 98 CZs for 2012-2017. The corresponding rental housing results are shown in Appendix Figure A8. Panel (c) shows time-to-build delays by completion year estimated from construction start-to-completion time microdata from the Census Survey of Construction. Region averages are estimated from regressions of individual project delays on region-by-year fixed effects. Estimates are broken out by U.S Census region.

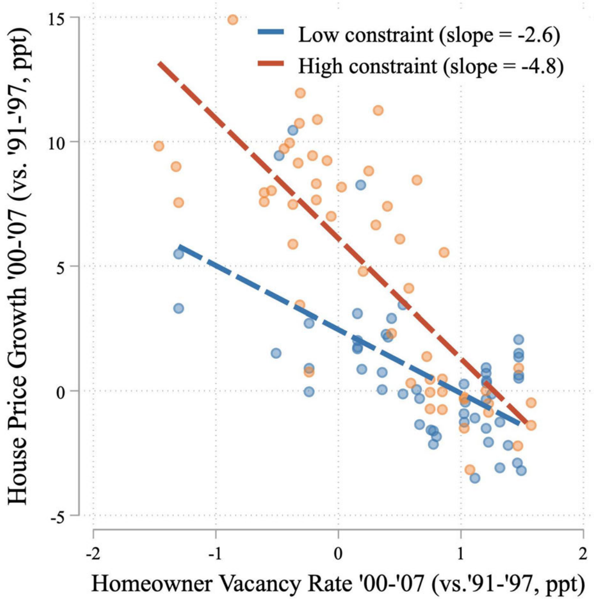

(a) Owner-occupied: 2000-2007

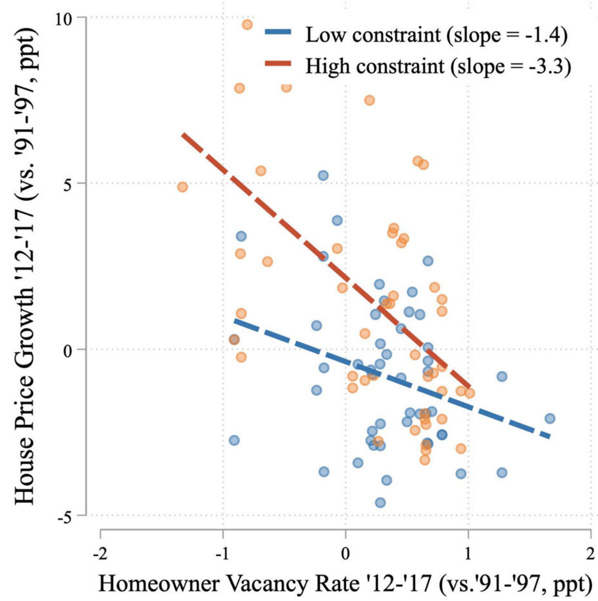

(b) Owner-occupied: 2012-2017

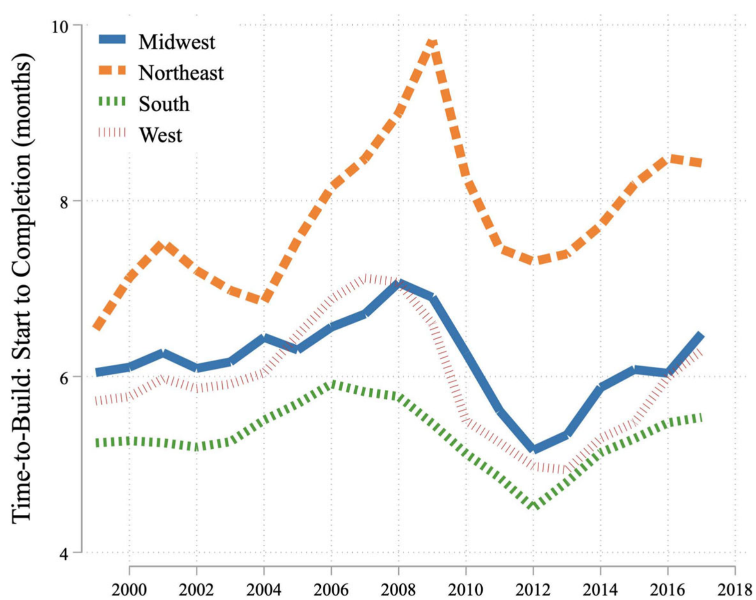

(c) Time-to-build delays by region

Figure 8: Migration flows by education group: evidence for displacement. The graphs in panels (a) and (b) plot gross out-migration of non-college workers (as % of group pop.) over gross in-migration rates of college-educated workers, at a CZ-level, pooling data from 2005-2017. Left panel shows all CZs, and right panel shows only CZs that fall into "high" or "low" housing supply constraint categories: High / low categories indicate land & regulatory constraints are *both* above / below median. Panel (c) shows the average difference in quality-adjusted house prices between origin and destination cities for movers of difference education groups.

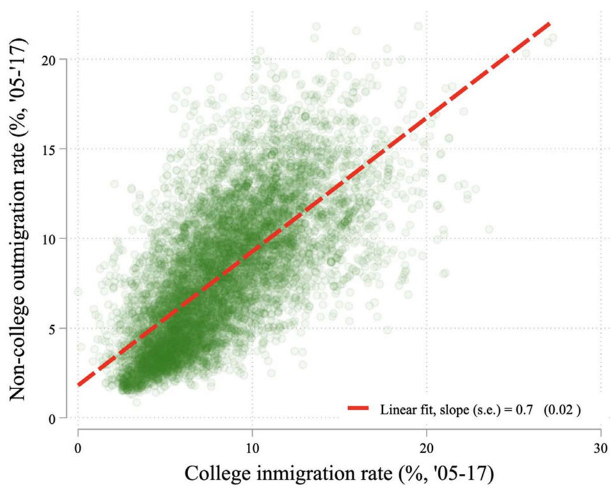

(a) College inflows & non-coll. outflows: All CZs

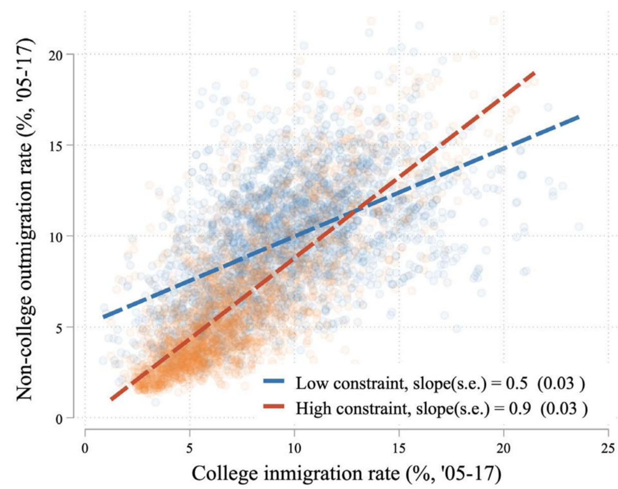

(b) College inflows & non-coll. outflows: high vs. low constraints

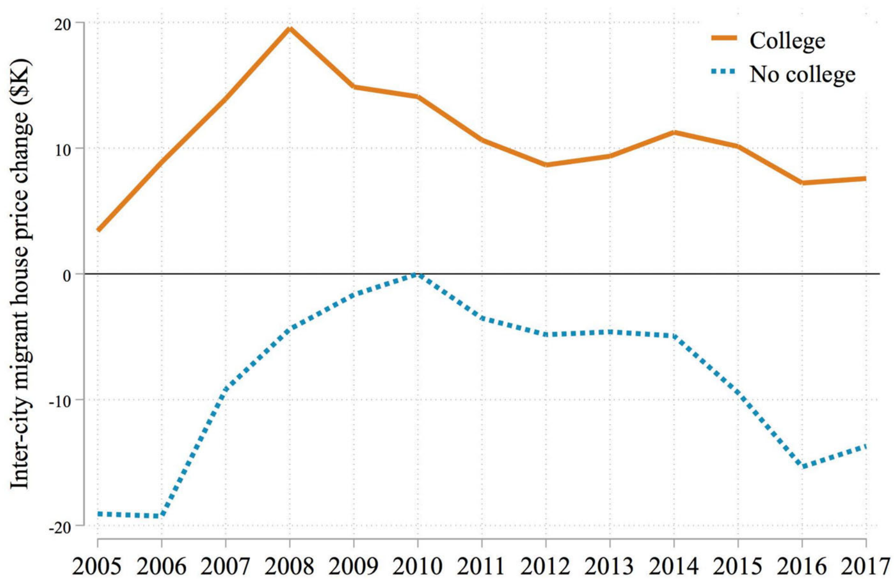

(c) House price change for movers

Figure 9: Predicted house price spillover examples 1995-2007. The graphs in Panels (a) and (b) plot the actual house price growth in 1995-2007 over the predicted spillover effects on house prices as a result of shocks to other cities consisting of: (1) Interest rate changes interacting with land constraints, and (2) Industry wage shock exposure due to industry employment shares, interacted with land constraints. The predicted spillover effects are computed as the predicted house price effects in other cities multiplied by the estimated network spillover coefficient from Column 6 of Table 4. The spillover effect variables has been residualized with regard to any direct effects on the focal city from the shock in question. Panels (c) and (d) show the effects on the house price beta for each city of the predicted spillover house price growth.

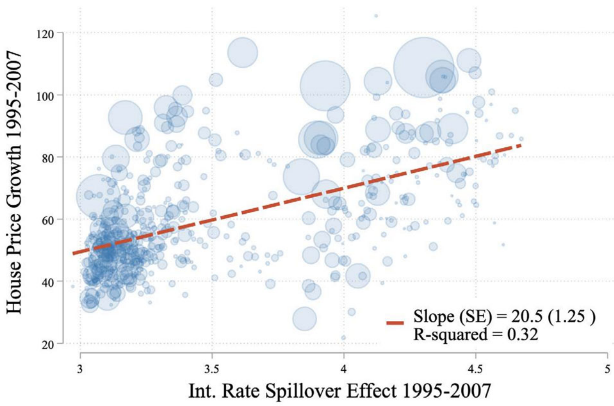

(a) House price effect: interest rate shock

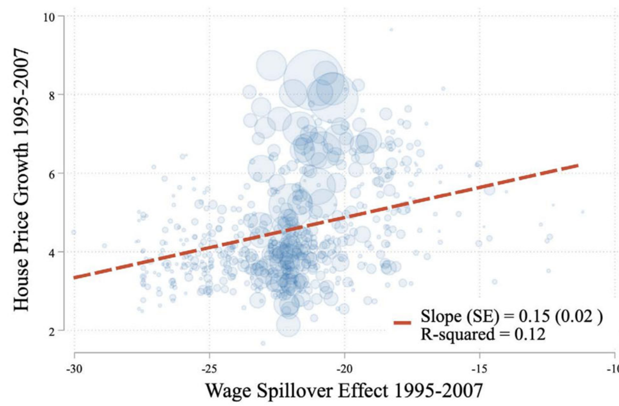

(b) House price effect: industry wage shock

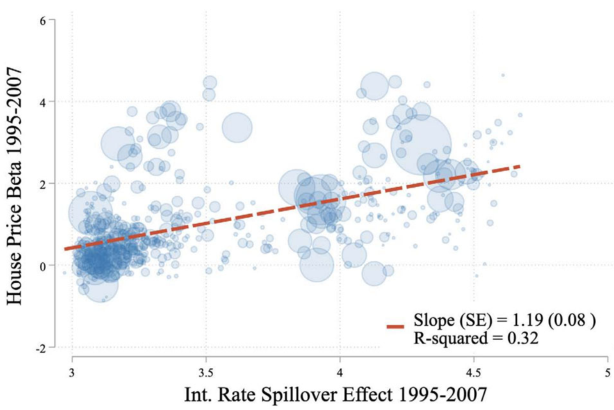

(c) Beta effect: interest rate shock

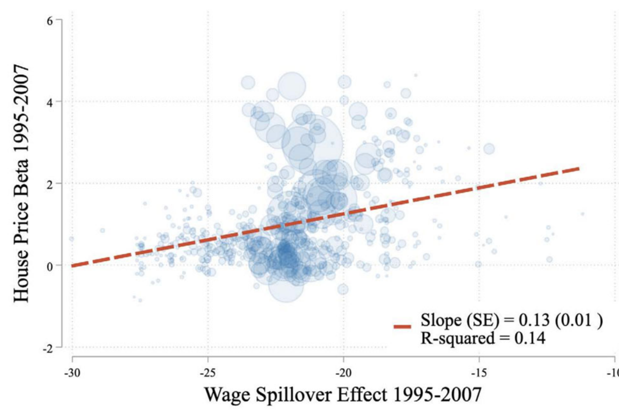

(d) Beta effect: industry wage shock

Figure 10: **Expected weighted house price correlations between CZs** The graph sorts CZs by their migration-weighted house price correlation, and then plots for each CZ the expected house price growth correlation 1990-2017 with all other CZs, weighting each CZ using the weights stated in the legend, which are (1) migration outflow shares (2) equal weights (3) inverse distance in miles (4) Facebook social connectedness weights, and (5) population weights (6) Trade flow weights. A higher value therefore indicates that the link captured by those weights is more strongly associated with a co-movement in house prices. The legend also shows the expected correlation using the weight, with numbers in parentheses being calculated only for the smaller sample of CZs where trade-flow-weighted expectations are available.

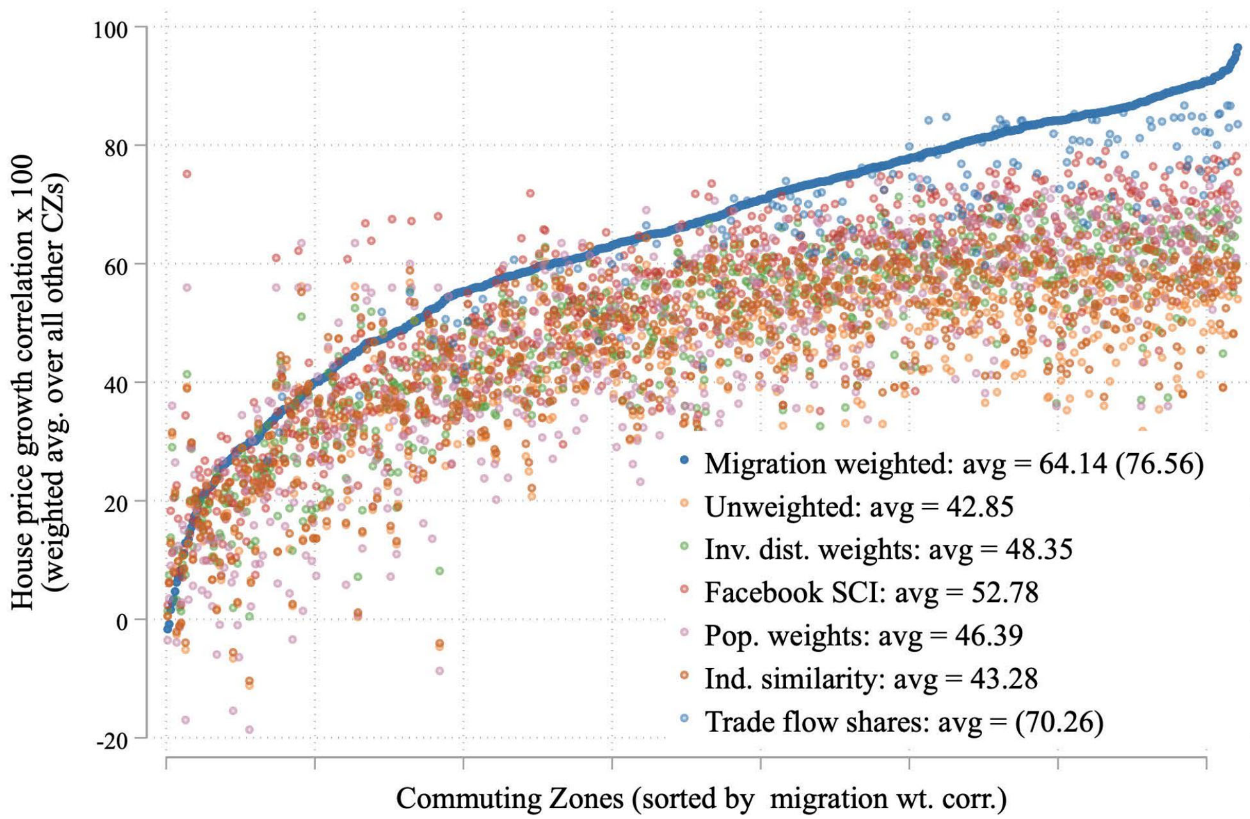

Figure 11: Sequence of events in the dynamic spatial equilibrium model for a group s worker starting period t in city i

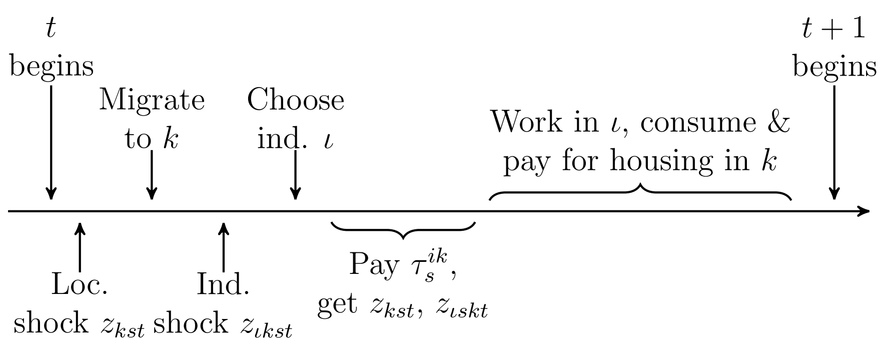

Figure 12: Illustration of "Renewal Action" identification. The figure shows how "renewal actions" allow us to identify the migration elasticity in the presence of future option value terms by comparing two different migration paths – and their associated utility – that end in the same location: Workers 1 and 2 end in the same location and therefore, by assuming that location choices are renewal states, have the same option values in the future. As a result, taking the difference in probabilities between these two paths will eliminate the unobserved future option value terms $E[V_{t+2}^{ks}]$.

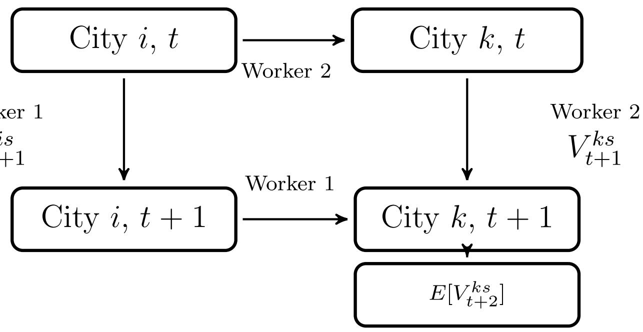

Figure 13: Model fit: historical imputation of college shares and log migration outflow shares in 2000. Left graph shows *actual* college shares by Adjusted CZ from the Year 2000 Census 5% IPUMS sample over college shares implied by the model-based imputation of 2000-2004 migration flows by education group. 512 Adjusted CZs are shown, corresponding to the continental U.S., excluding New Orleans. Right panel shows the model fit for the log of total outmigration shares for all destination-origin pairs. Vertical axis shows the actual value in the IRS data, while the horizontal axis is the value obtained by adding up the model-imputed flows by education group for each destination-origin pair for the year 2000.

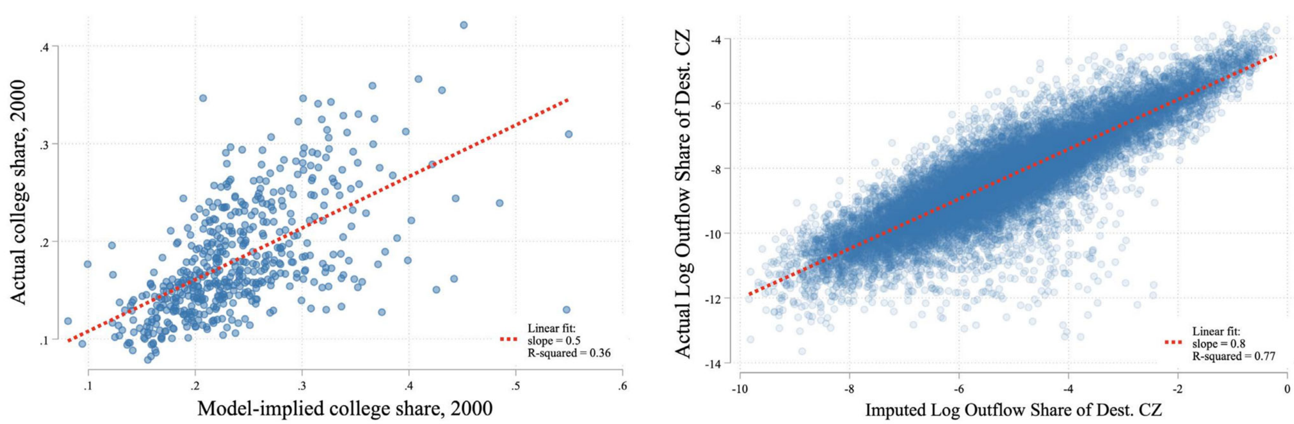

(a) College shares in 2000 **(b) Migration outflow shares in 2000**

Figure 14: Cumulative nominal industry wage shocks Graphs show the total change in nominal wages implied by national industry wage trends where the left (right) panels assume that each NAICS 2-digit industry experiences seven (six) years of its national nominal wage trend over 2000-2007 (2012-2017). Both panels show all 512 continental Adjusted CZs (mapped onto 1990 CZ boundaries), excl. New Orleans.

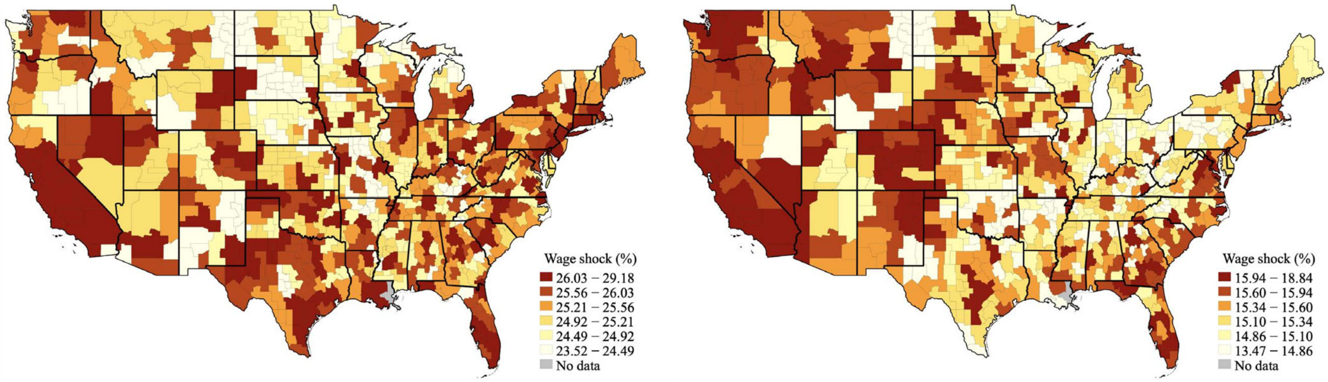

(a) 2000-2007 **(b) 2012-2017**

Figure 15: Comparison of actual house price growth to simulated wage shock effect. The graphs plot actual house price growth for 2000-2007 (Panel A) and 2012-2017 (Panel B) for 512 Adjusted CZs relative to the house price growth relative to baseline in the same period in the model for a stationary steady-state transition path where the industry wage shocks are applied.

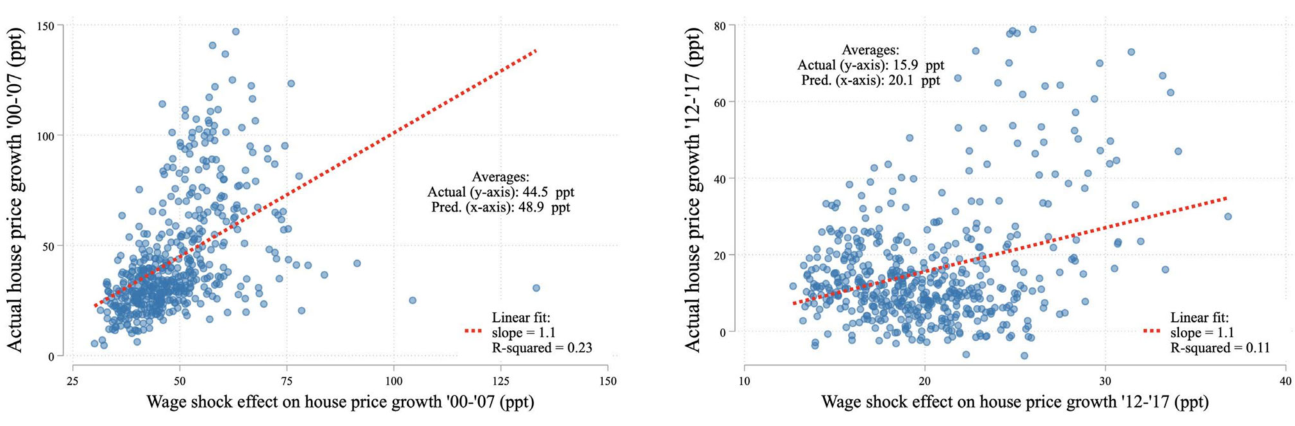

(a) 2000-2007 **(b) 2012-2017**

Figure 16: Counterfactual 1: Effect of migration cost changes. The graphs in all panels show smoothed kernel density plots, with an Epanechnikov kernel with bandwidth: 1 in (a), and 2 in (b) and 0.5 in (c) and (d). Panels (a) and (b) show the distribution of house price growth in response to wage shocks in Adjusted CZs. In order, the orange curve shows the decreased migration cost (i.e. increased mobility) scenario; the blue curve shows the wage shock effects with baseline mobility; the red curve shows the distribution under the scenario with prohibitive inter-state migration costs; and the green curve shows the no mobility scenario with constant populations. Panels (c) and (d) show the distribution of betas for the annual house price growth effect in each period, computed by regressing each city's annual growth series in each counterfactual on the series of leave-one-out average growth in each year.

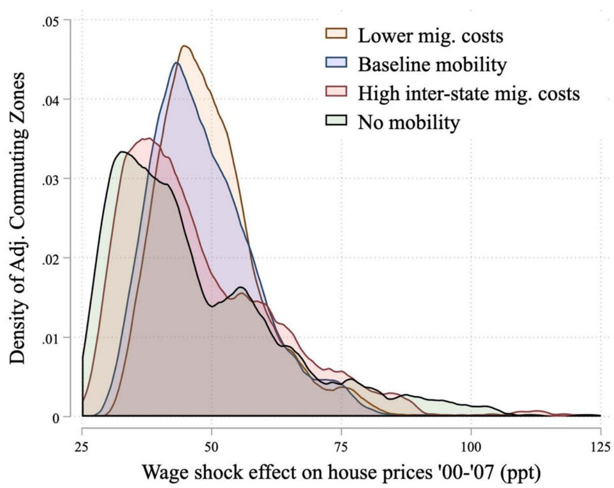

(a) House price growth effect: 2000-2007

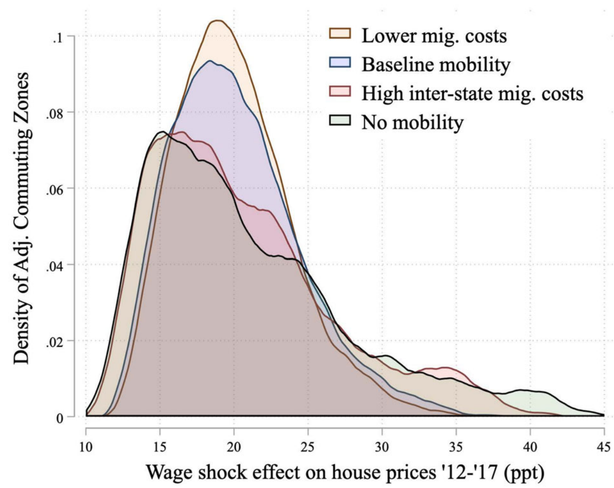

(b) House price growth effect: 2012-2017

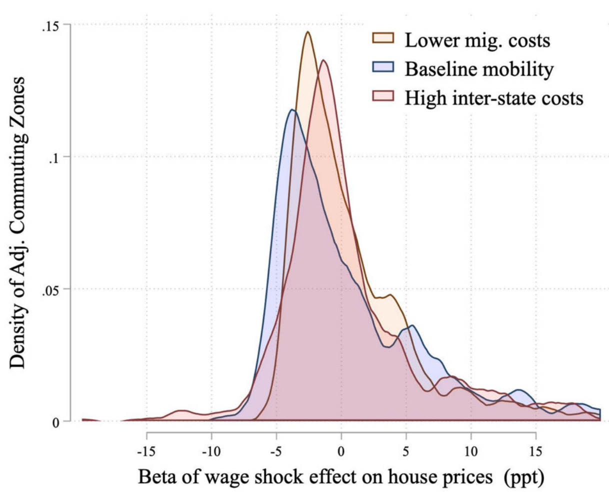

(c) House price effect beta: 2000-2007

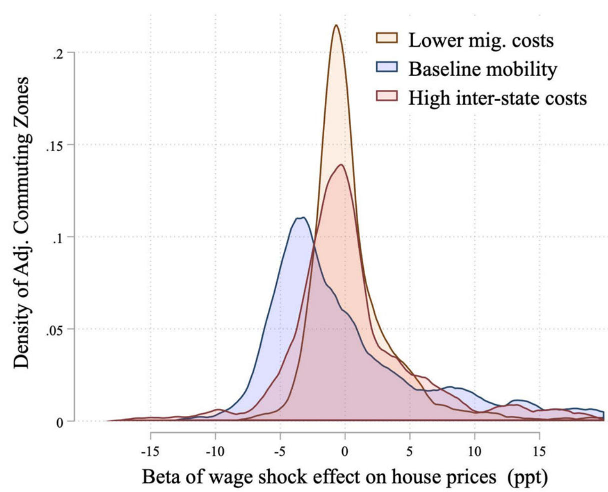

(d) House price effect beta: 2012-2017

Figure 17: Counterfactual 2: reductions in house price elasticity. Maps in Panels (a) and (b) show the reduction in the inverse housing supply elasticity implied by reducing constraints down to the 75th pctl. of inverse elasticity, or to the median. Panels (c) and (d) show smoothed kernel density plots (Epanechnikov kernel with bandwidths 1 (a), 2 (b) and 0.5 (c and d)) of the distribution of house price growth in response to wage shocks in Adjusted CZs. Panels (e) and (f) show the distribution of betas for the annual house price growth effect in each period, computed by regressing each city's annual growth series in each counterfactual on the series of leave-one-out average growth in each year.

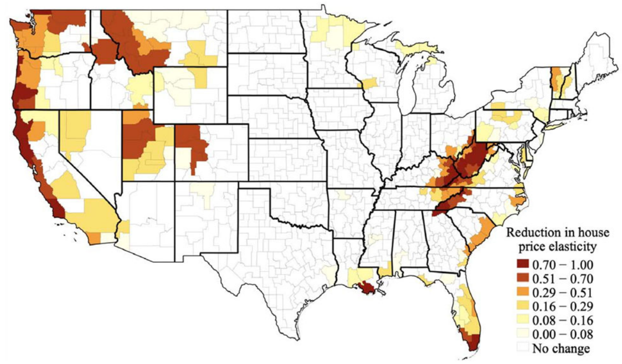

(a) Reduction to P75

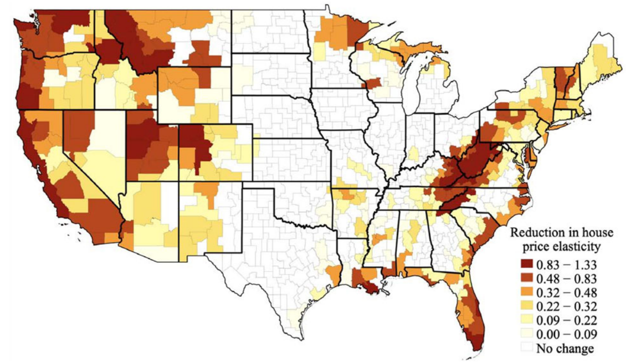

(b) Reduction to Median

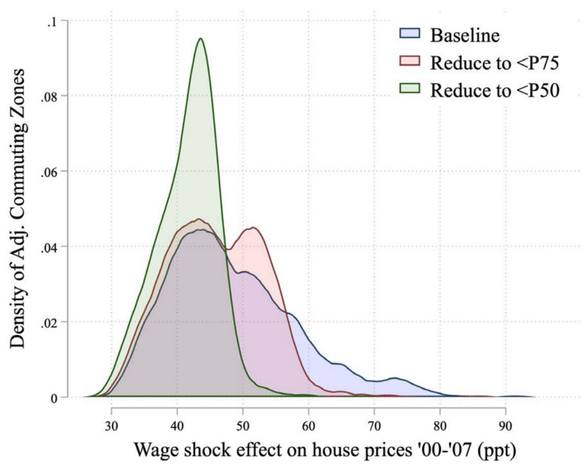

(c) Total house price effect: 2000-2007

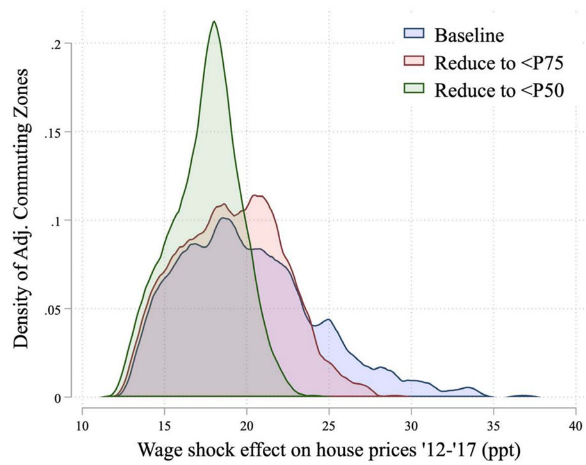

(d) Total house price effect: 2012-2017

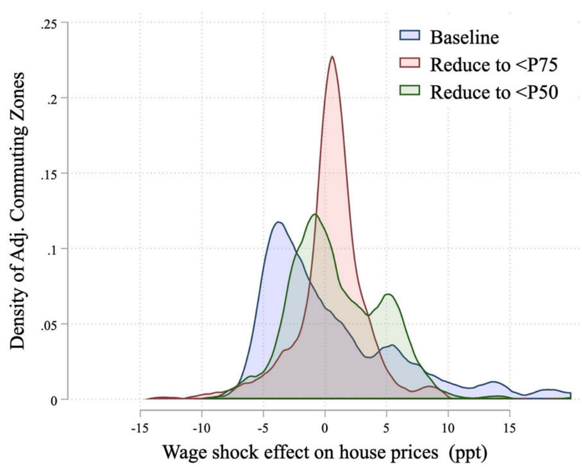

(e) House price effect beta: 2000-2007

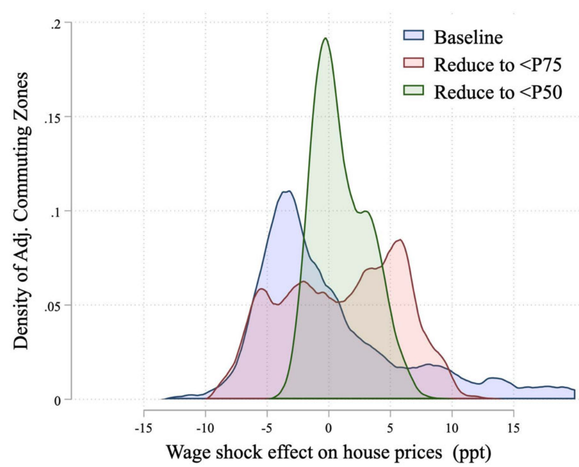

(f) House price effect beta: 2012-2017

C Appendix Figures

Figure A1: House price growth and inter-city migration trends. Panel (a) shows population-weighted averages of gross inmigration rates in U.S. CZs from domestic origins and CZ real house price growth over time. Gross migration rates for 2013-2017 have been smoothed due to a data issue in the IRS gross migration data for that period, which generates excess volatility. Panel (b) plots the same gross migration rates together with the spread between the 90th and the 10th percentile (population-weighted) of cities in real house price growth in each year.

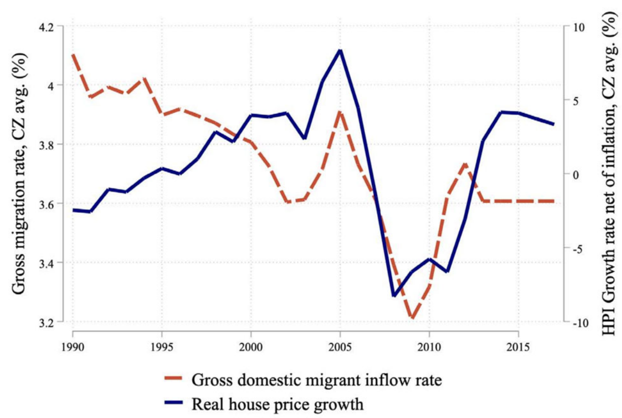

(a) Gross migration and house price growth avg.

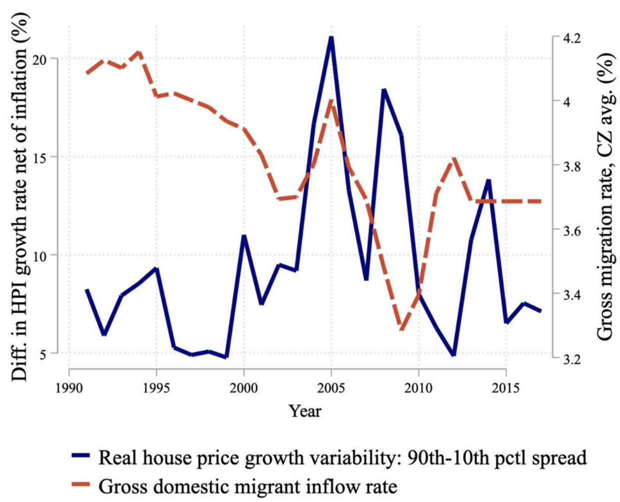

(b) Gross mig. and spread in house price growth

Figure A2: Migration inflow network examples. Shading in the map indicates % of continental U.S. inflows to the city coming from the CZ over 2000-2007 period in IRS migration data.

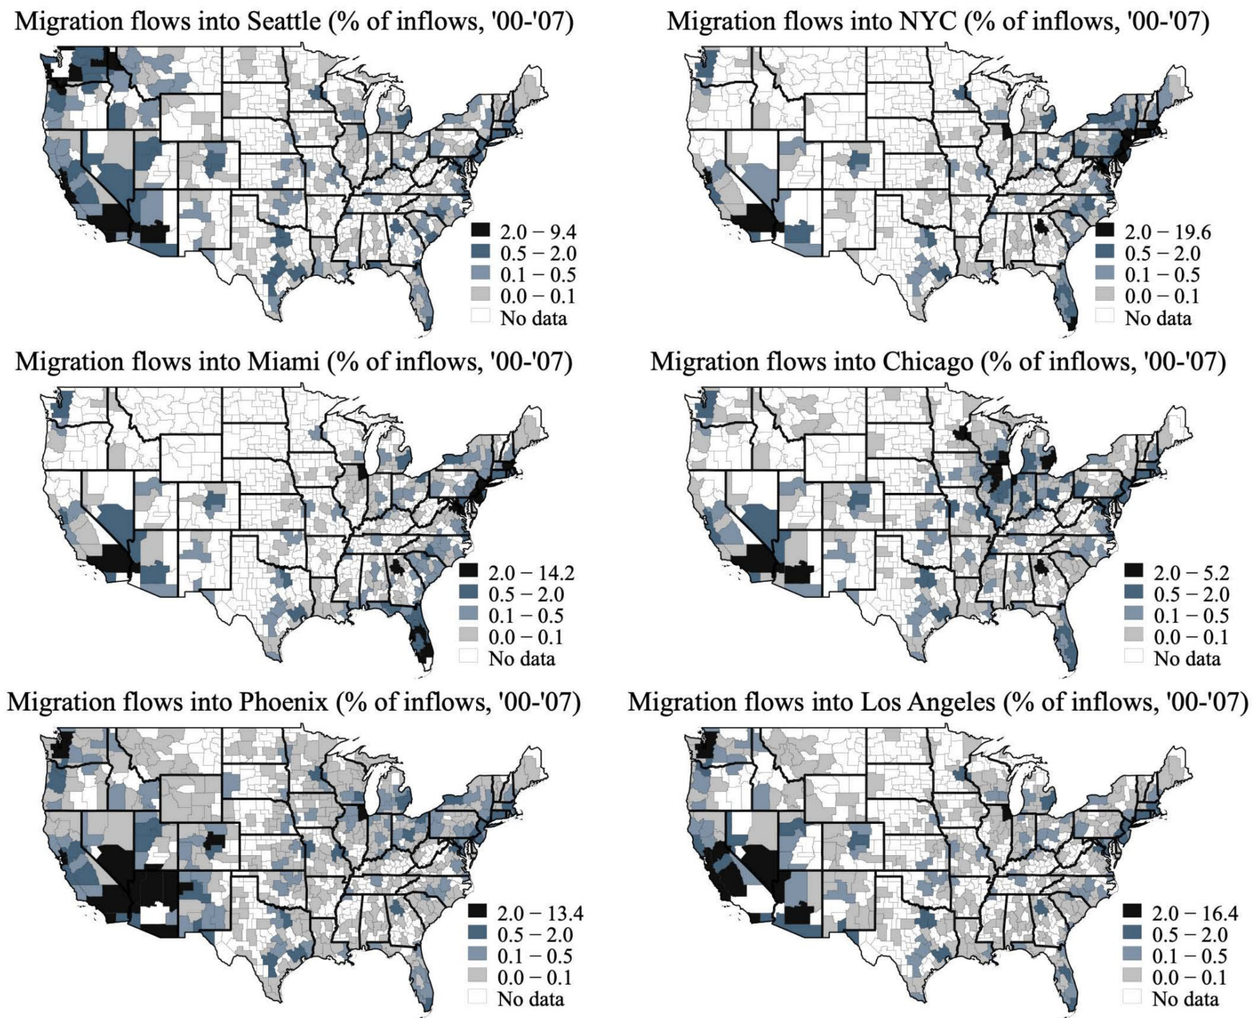

Figure A3: **First stage instruments and network house price growth.** Graphs show binned scatter plots of the pairwise first-stage relationship between network house price growth and the network house price growth instruments. All variables are residualized with regard to city and year FEs, regional trends, migration access, local Bartik shock, local Bartik shock × land unavailability interaction, local log wage growth and network wage growth. Data pools years 1991-2017 for 586 CZs and includes migration network with > 150 mi. distance from each city. The left panel shows the network instrument NWP_{it}^{Bx} that is based on the interaction of Bartik shocks with land unavailability, and the right panel shows the first stage slope for NWP_{it}^{B} , which only uses Bartik shocks.

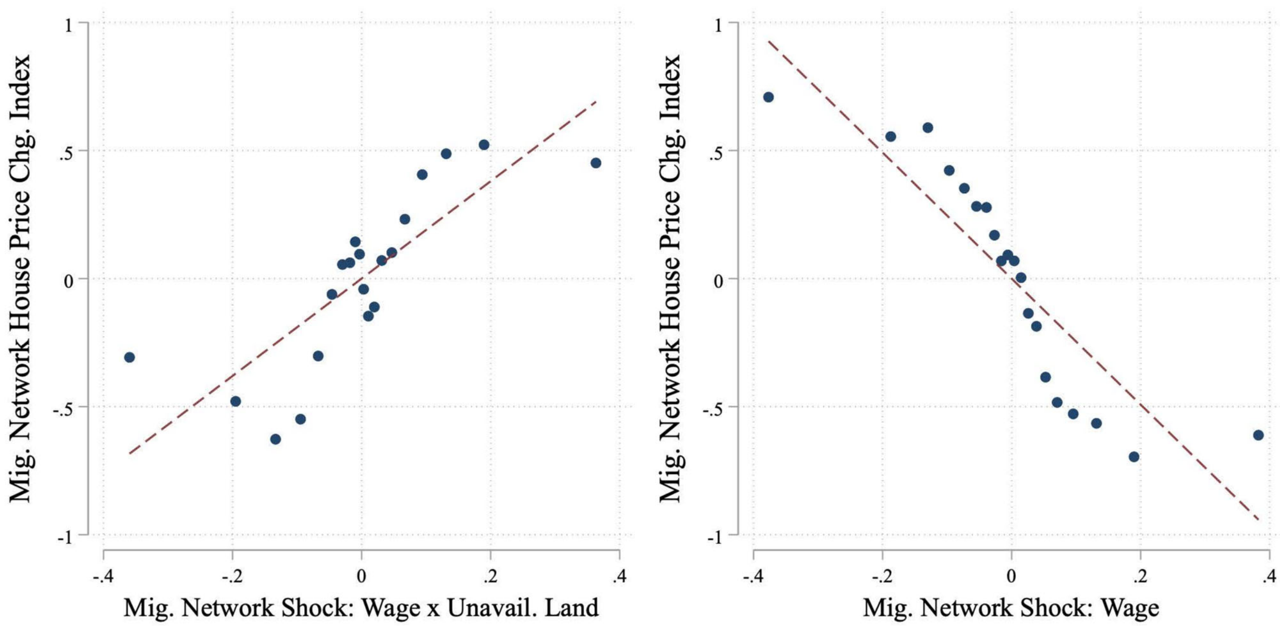

Figure A4: OLS dynamic network spillover effects on house prices Graphs show local projection coefficients corresponding to the correlation of period t network house price growth with period $t-1+h$ outcome variables of the form

$$\sum_{k=1}^{h} \Delta \ln P_{i,t-1+h} = \alpha_i + \alpha_t + \tilde{\eta}^{nw}_{cum} \Delta \mathcal{P}^{NW}_{it} + \beta'_{cum} \Gamma_{it} + \tilde{\xi}^{P,cum}_{i,t-1+h}.$$

Control variables in all regressions include year & CZ FEs, regional trend FEs, migration access, local log avg. salary growth, network log salary growth and local Bartik productivity shocks. Regressions in panels a, b, e, and f also include an interaction between local land share unavailable for construction and the productivity shocks, while panels c an d additionally control for interactions of migration access, as well as local and network log salary growth, with the local land share unavailable for construction. The extrapolation regressions in panels e and f include one lag of local and network house price growth in addition to the contemporaneous network house price growth for which the coefficient is displayed. Estimation uses data for 1991-2017. Dashed lines show 95% CI based on std. errors clustered at the CZ level.

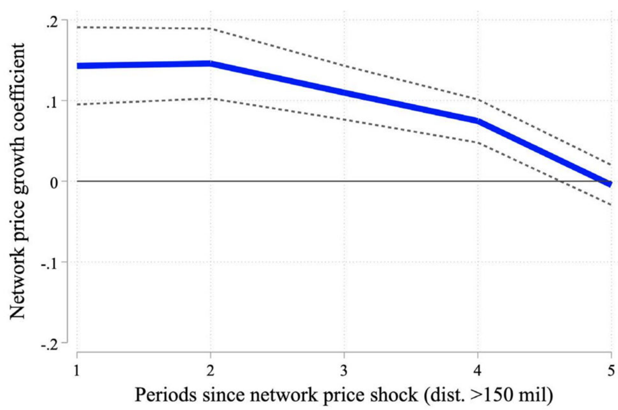

(a) Log House Price Growth: dist. > 150 mi.

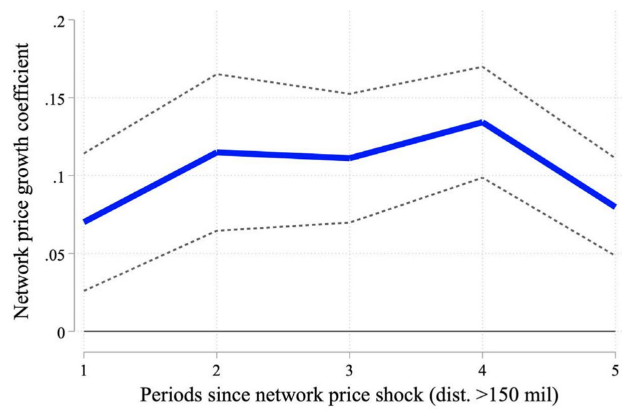

(b) Extrapolation: $\Delta \mathcal{P}^{NW}_{it}$ coeff. (dist. > 150 mi.)

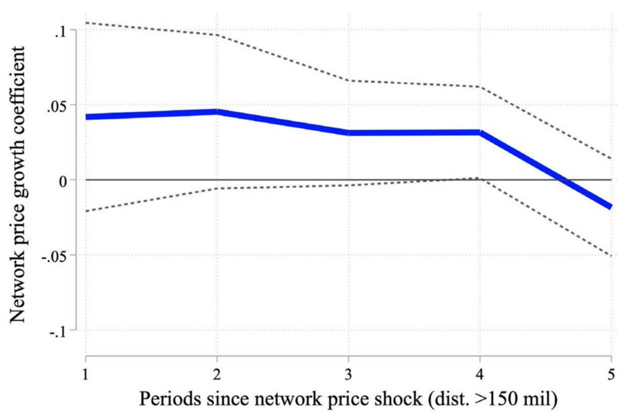

(c) Heterogen.: $\Delta \mathcal{P}^{NW}_{it}$ coeff. (dist. > 150 mi.)

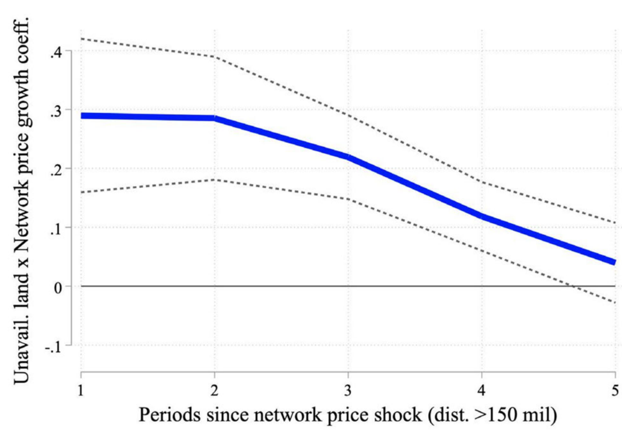

(d) Heterogen.: $\Delta \mathcal{P}^{NW}_{it} \times x^{\text{land}}_i$ coeff. (dist. > 150 mi.)

Figure A5: Dynamic effects: other outcomes - OLS Graphs show local OLS coefficients corresponding to the effect of period t network house price growth with period $t-1+h$ outcome variables of the form

$$Y_{i,t-1+h} = \alpha_i + \alpha_t + \tilde{\eta}_h^{nw} \Delta \mathcal{P}_{it}^{NW} + \beta' \Gamma_{it} + \tilde{\xi}_{i,t-1+h}^{P}.$$

Baseline control variables Γ_{it} in all regressions include year & CZ FEs, regional trend FEs, migration access, local log avg. salary growth, network log salary growth and local Bartik productivity shocks, as well as local productivity shocks interacted with local land share unavailable for construction. Baseline instruments for network house price growth in all regressions consist of the migration-weighted network averages of productivity shocks, and of migration-weighted productivity shocks interacted with local land share unavailable for construction. Estimation uses data for 1991-2017. Dashed lines show 95% CI based on std. errors clustered at the CZ level.

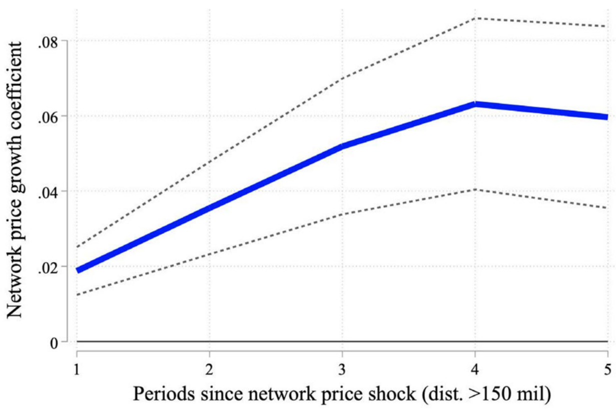

(a) OLS: Cumul. Pop Impact of Net Mig. (log pts)

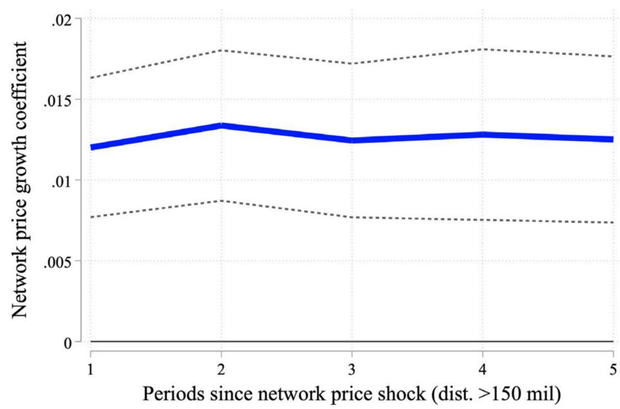

(b) OLS: Log Housing Unit Permits

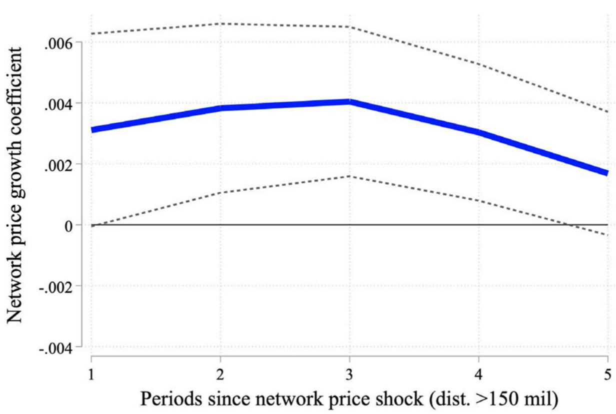

(c) OLS: Log Mortgage Purchase Loans

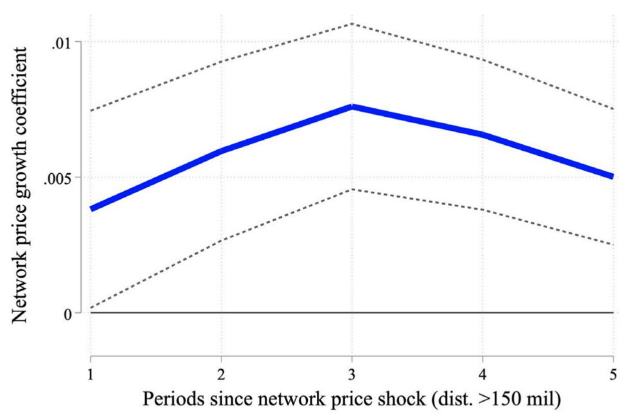

(d) OLS: Log Mortgage Lending Volume (in USD)

Figure A6: House price growth and net migration. The graph plots CZ average annual house price growth and average net migration rates during different time periods. All period averages are net of their averages during 1991-1999. For example, the fitted line for the 2000-2007 boom period corresponds to the regression

$$\overline{\Delta \ln P}_{\text{Boom},i} - \overline{\Delta \ln P}_{\text{‘}91\text{-‘}97,i} = \alpha + \beta(\overline{NetMig\%}_{\text{Boom},i} - \overline{NetMig\%}_{\text{‘}91\text{-‘}97,i}) + \epsilon_i$$

Graph and fitted line include all 559 CZs with population > 30K, but estimate line slopes represent CZs of all population sizes, weighted by their population and excluding New Orleans.

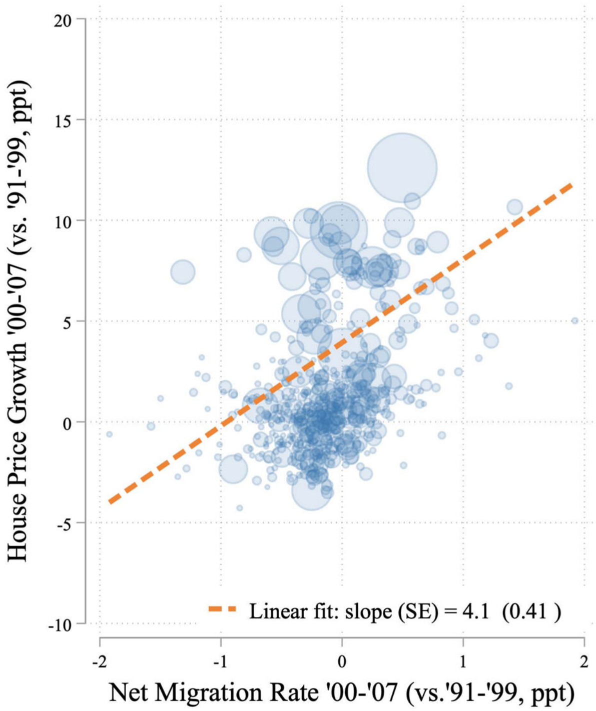

(a) Boom: 2000-2007

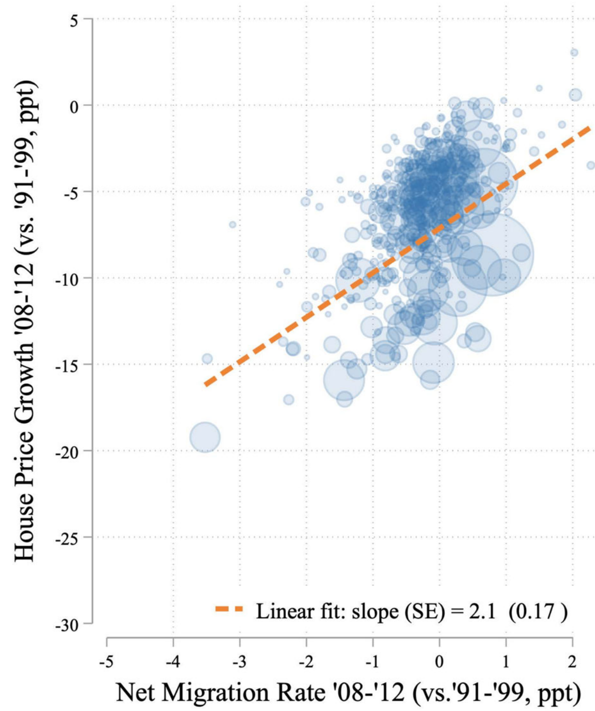

(b) Bust: 2008-2012

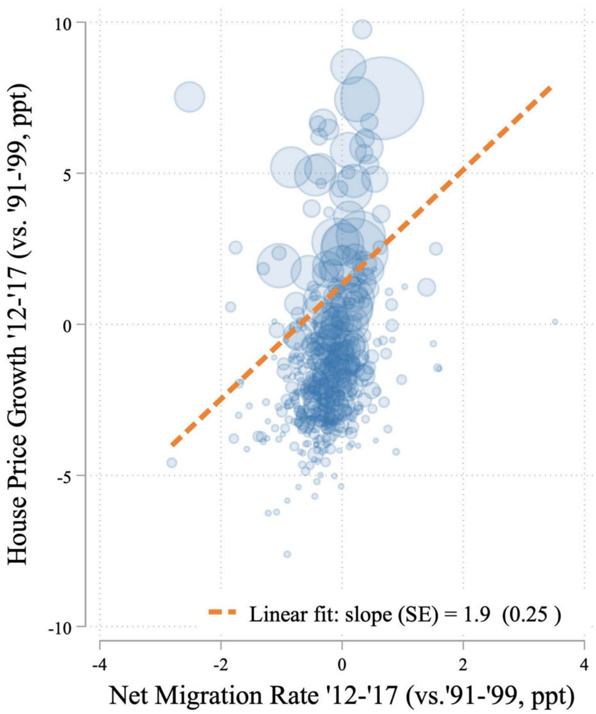

(c) Boom: 2012-2017

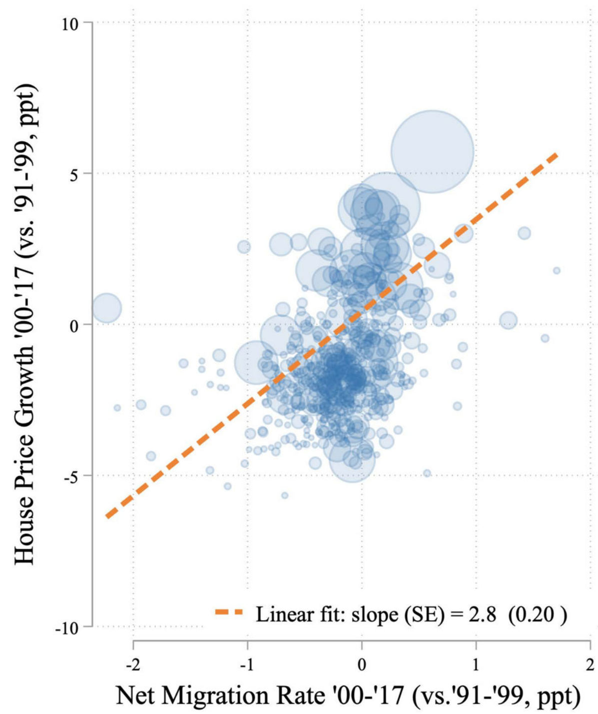

(d) 2000-2017

Figure A7: Expected weighted house price correlations between CZs The graph sorts CZs by their migration-weighted house price correlation, and then plots for each CZ the expected house price growth correlation 1990-2017 with all other CZs, weighting each CZ using the weights stated in the legend, which are (1) migration outflow shares (2) equal weights (3) inverse distance in miles (4) Facebook social connectedness weights, and (5) population weights (6) Trade flow weights. A higher value therefore indicates that the link captured by those weights is more strongly associated with a co-movement in house prices. The legend also shows the expected correlation using the weight, with numbers in parentheses being calculated only for the smaller sample of CZs where trade-flow-weighted expectations are available. Each panel excludes CZs at less than the stated distance from the computation of a CZ's average housing correlation with each weight.

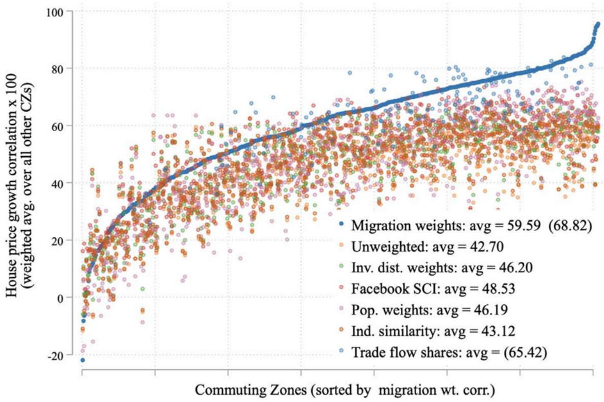

(a) Distance > 50 mi.

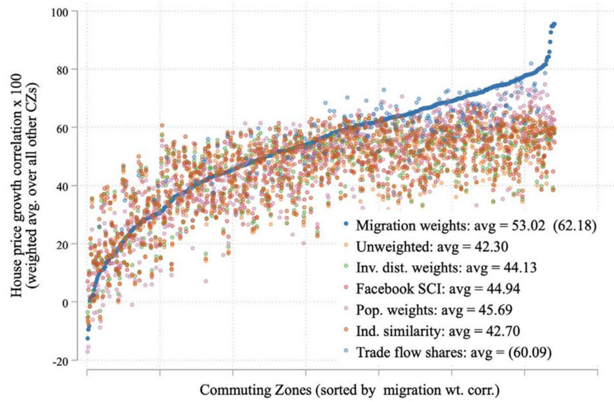

(b) Distance > 150 mi.

Figure A8: House price growth and vacancy rates: rental housing. The graphs plot CZ average annual house price growth during the housing booms of 2000-2007 and 2012-2017 over the change in average vacancy rates for rental housing. All averages are computed as differences relative to the 1991-1997 average. The graphs and fitted lines include 104 CZs for 2000-2007 and 98 CZs for 2012-2017.

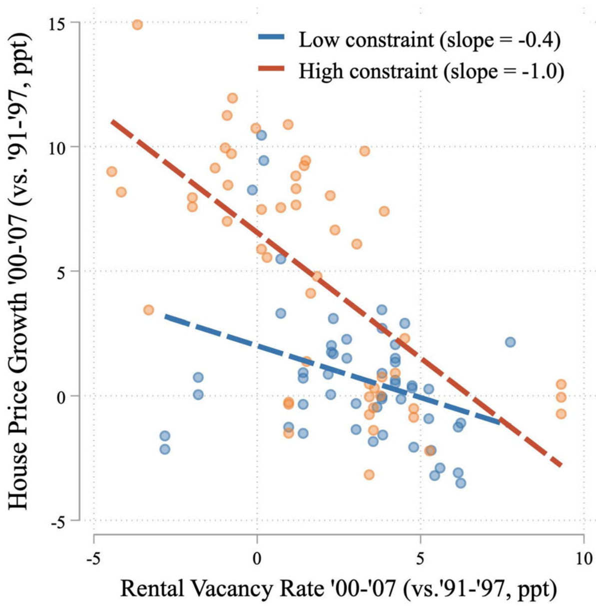

(a) Rental: 2000-2007

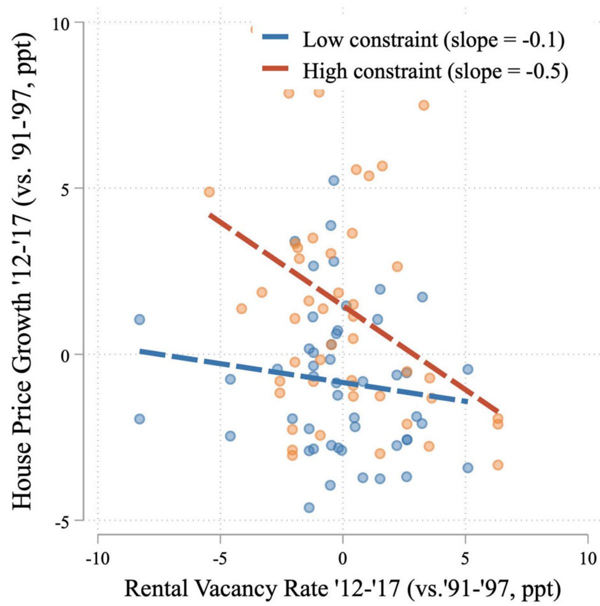

(b) Rental: 2012-2017

Figure A9: In- and outmigration by education and supply constraints. The graph plots gross out-migration of non-college workers (as % of group pop.) over gross in-migration rates of college-educated workers, at a CZ-level, pooling data from 2005-2017. CZs are divided by housing supply constraints: "High" constraints (low) indicate land & regulatory constraints are both (not) above median.

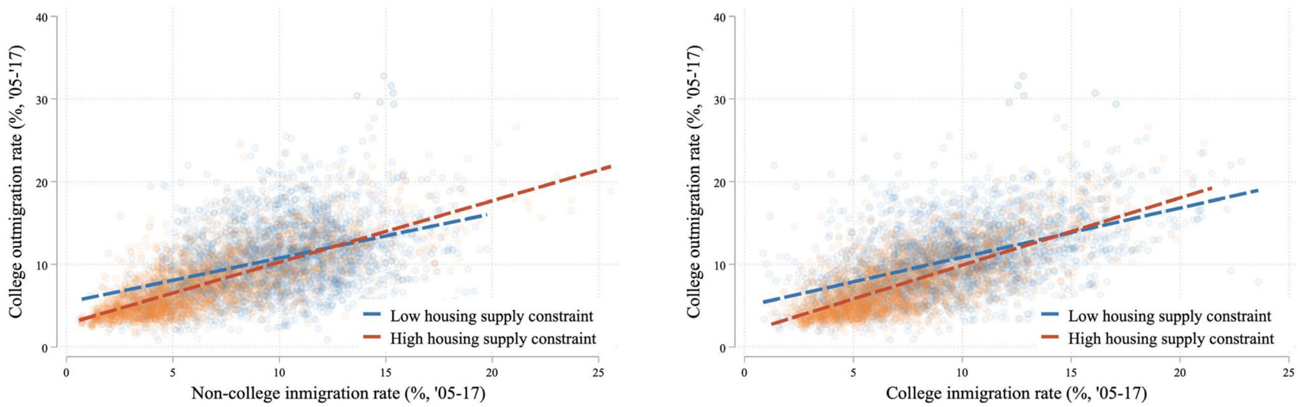

(a) Non-college inflows & college outflows **(b) College inflows & college outflows**

Figure A10: Counterfactual 1: Effect of migration cost changes on migration distances. The graphs in all panels show smoothed kernel density plots, with an Epanechnikov kernel with bandwidth 0.1. The plots show the distribution of year 2012 inter-city migration flows over the log distances between Adjusted CZs. The blue curve shows the baseline of observed flows and the red curve the distribution under the scenario with prohibitively high inter-state migration costs. College worker flows are shown in Panel (a), and non-college worker flows in Panel (b).

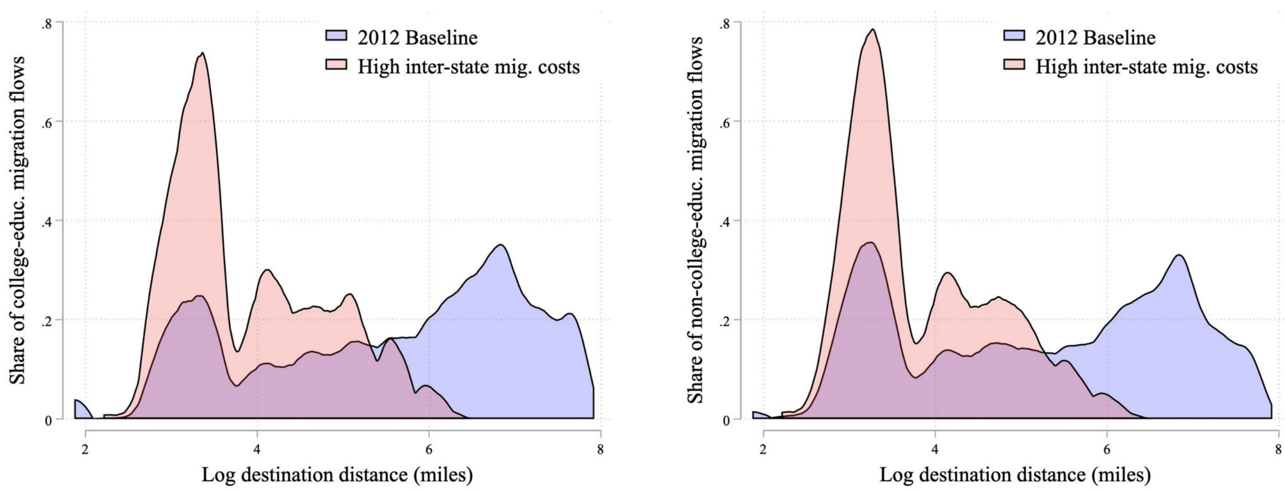

(a) College flows **(b) Non-college flows**

Figure A11: Migration spillover example – Boston, MA, and Portland, ME. The graph shows net migration flows (IRS) between Boston and Portland, as well as house price growth in both cities.

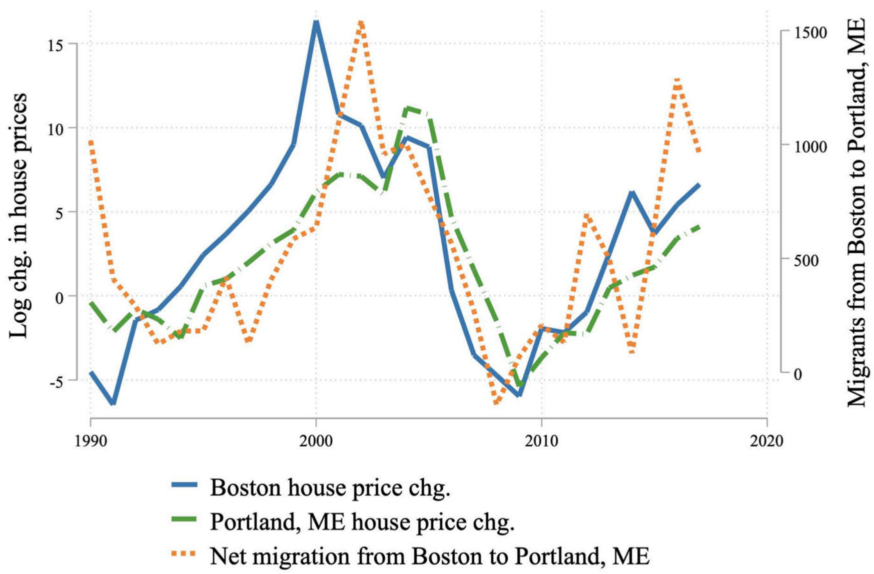

D Appendix Tables

Table A1: Summary statistics. The table below shows summary statistics for the key variables used in the different estimations. The three panels correspond to different samples: Panel A shows the data available for 1990-2017 for the continental U.S. CZs (under the 1990 boundary definition), of which there are 721 (excl. New Orleans). Panel B describes the data for the 2000-2017 sample of 512 Adjusted CZs in the continental U.S. with at least 50K residents, which is used in the city-level quantitative parameter estimations. Panel C shows the summary statistics for the 2-digit NAICS industry-by-Adjusted-CZ panel for the years 2000-2017 used in the labor demand and industry choice parameter estimations. Panel D describes the key city pair variables used in the location choice estimation.

Sample	Mean	SD	Median	Min.	Max.	Obs.
Panel A: Reduced-form CZs 1990-2017						
Population ('000s, IRS)	345.82	948.80	89.95	1	17360	20,188
Log chg. in house prices (FHFA)	2.92	4.50	2.97	-41	44	20,188
Net domestic migration (% of pop., IRS)	0.05	6.09	-0.04	-50	851	20,188
Gross domestic in-migration (% of pop., IRS)	4.70	6.26	4.37	0	858	20,188
Gross domestic out-migration (% of pop., IRS)	4.65	1.64	4.37	0	52	20,188
Bartik wage shift-share shock (1990 wts., 3-dig.)	3.08	1.27	3.07	-11	19	19,467
Construction permits ('000 housing units)	1.86	5.05	0.28	0	90	19,681
Purch. mortgage originations ('000s, annual)	5.35	16.11	0.79	0	429	20,055
Purch. mortgage orig. vol. (Mil. USD, annual)	922.42	3818.25	77.95	0	137882	20,055
Unavailable land share (%, Lutz & Sand, 2019)	26.01	21.67	19.88	0	92	19,170
Network house price chg. exposure	3.00	3.98	3.39	-24	41	19,170
Network productivity shock	3.18	1.28	3.08	-1	7	19,170
Network land share× productivity shock	0.77	0.59	0.63	-0	4	19,170
Panel B: Adjusted CZs 2000-2017						
Population ('000s, IRS)	465.91	1037.90	156.30	39.0	15086.9	9,216
Log chg. in house prices (log pts, FHFA)	2.34	4.90	2.42	-39.8	28.8	9,216
Bartik wage shift-share shock (2000 wts., 3-dig.)	2.80	1.09	2.90	-1.8	6.7	9,216
Annual wages: city avg. ('000 USD, QCEW)	33.95	8.33	32.83	17.7	113.2	9,216
Annual wages: college('000 USD, QCEW)	36.78	10.36	35.12	17.6	143.8	9,216
Annual wages: non-college ('000 USD, QCEW)	33.20	7.52	32.36	17.6	96.1	9,216
Unavailable land share (%, Lutz & Sand, 2019)	27.09	20.49	22.64	0.3	90.5	9,216
College share of pop. (%, ACS)	21.86	6.72	20.55	7.4	50.5	6,656
Panel C: Industry-by-Adj.-CZ 2000-2017						
Industry employment ('000s, QCEW)	10.89	35.13	1.81	0	1096	174,116
Log chg. in industry employment (log pts)	1.15	26.15	0.77	-540	639	166,882
Avg. annual industry wages ('000 USD, QCEW)	36.93	20.07	33.38	2	1089	168,364
Log chg. in industry wages (log pts)	2.99	8.44	2.88	-256	323	166,882
Panel D: Adjusted CZ Pairs 2005-2017						
College migration share $i \to k$ (%, smoothed)	0.013	0.080	0.001	0.000	7.425	3,401,216
Non-coll. migration share $i \to k$ (%, smoothed)	0.007	0.045	0.001	0.000	4.138	3,401,216

Table A2: Reduced-form house price spillover effects - by network distance

Dependent variable:	CZ Log House price growth$_{it}$			
Network min. distance	0 mi.	50 mi.	100 mi.	150 mi.
	(1)	(2)	(3)	(4)
Panel A: OLS				
Δ Network HP$_t$	0.86***	0.57***	0.31***	0.14***
	(0.04)	(0.04)	(0.03)	(0.02)
Panel B: IV				
Δ Network HP$_t$	0.69***	0.49***	0.35***	0.23***
	(0.10)	(0.08)	(0.07)	(0.06)
Observations	19,170	18,063	16,794	15,822
1st-stage F-stat.	78	115	94	115
Year FE	✓	✓	✓	✓
CZ FE	✓	✓	✓	✓
Regional trend FEs	✓	✓	✓	✓
Migration Access$_{i,t-1}$	✓	✓	✓	✓
ΔW_{it}^{NW}	✓	✓	✓	✓
Wage shock$_{it}$	✓	✓	✓	✓
Wage shock$_{it}$ $\times$ Unavail. Land$_i$	✓	✓	✓	✓

Heteroskedasticity-robust standard errors clustered at the CZ level in parentheses: * $p<0.10$, ** $p<0.05$, *** $p<0.01$. Includes data from 586-709 CZs for 1991-2017. Migration network by distance includes all CZs for which the center of no component county is closer than the stated number of miles to the center of any focal city component county. Network house prices are the weighted sum of other CZs' house price changes. Migration access is the migration-weighted sum of city populations. Both measures use average 1990-1995 migration flows to compute migration weights. The "Wage shocks" are Bartik shocks computed as a weighted average of national leave-one-out wage growth by industry, with weights given by local 3-digit NAICS industry shares in 1990.

Table A3: Reduced-form house price spillovers: long-run effects

Dependent variable:	Δ Log House price$_{i,t+4}$			Δ Log Pop.$_{\cdot i,t+4}$	Log Permits$_{i,t+4}$
	(1)	(2)	(3)	(4)	(5)
Δ Network HP$_t$	0.626***	0.431**	0.633*	0.186***	0.027***
	(0.167)	(0.188)	(0.324)	(0.046)	(0.009)
ΔHP_t × Unavail. land		0.944***			
		(0.197)			
House price growth$_{t-1}$			2.938***		
			(0.508)		
Δ Network HP$_{t-1}$			-0.700**		
			(0.350)		
Observations	13,478	13,478	12,306	13,478	13,052
Year FE	✓	✓	✓	✓	✓
CZ FE	✓	✓	✓	✓	✓
Regional trend FEs	✓	✓	✓	✓	✓
Migration Access$_{i,t-1}$	✓	✓	✓	✓	✓
ΔW_{it}^{NW}	✓	✓	✓	✓	✓
Wage shock$_{it}$	✓	✓	✓	✓	✓
Wage shock$_{it}$ × Unavail. Land$_i$	✓		✓	✓	✓

Heteroskedasticity-robust standard errors clustered at the CZ level in parentheses: * p<0.10, ** p<0.05, *** p<0.01. Includes data from 586 CZs for 1991-2017. Migration network by distance includes all CZs for which the center of no component county is closer than 150 miles to the center of any focal city component county. Network house prices are the weighted sum of other CZs' house price changes. Migration access is the migration-weighted sum of city populations. Both measures use average 1990-1995 migration flows to compute migration weights. The "Wage shocks" are Bartik shocks computed as a weighted average of national leave-one-out wage growth by industry, with weights given by local 3-digit NAICS industry shares in 1990. All estimates shown for house price and population are cumulative total effects of a period t shock on period $t+4$ log house prices and population (the latter as a result of domestic net migration). Permit effects are long-run effects on the annual *flow* of permits.

Table A4: Location choice parameters: Additional OLS results

Education group:	College				No College			
Model:	No Coll. Share (1)	Amen. Index (2)	CZ Inc. per cap. (3)	Static dep. var. (4)	No Coll. Share (5)	Amen. Index (6)	CZ Inc. per cap. (7)	Static dep. var. (8)
Δ Log College Real Wage diff.	0.03	0.03		0.13*				
	(0.10)	(0.10)		(0.07)				
Δ Log Non-Coll. Real Wage diff.					0.00	-0.02		0.21***
					(0.07)	(0.07)		(0.05)
Δ College share diff.		2.20***	2.20***	1.68***		-0.58***	-0.59***	-0.42***
		(0.31)	(0.31)	(0.17)		(0.18)	(0.18)	(0.10)
Δ Amenities index diff.		-0.02				-0.07***		
		(0.02)				(0.01)		
Δ Log CZ Real IRS Inc. diff.			0.26*				0.15*	
			(0.15)				(0.08)	
Observations	2,877,952	2,877,952	2,877,952	3,139,584	2,877,952	2,877,952	2,877,952	3,139,584
Year FE	✓	✓	✓	✓	✓	✓	✓	✓

Heteroskedasticity-robust standard errors clustered at the CZ pair level in parentheses: * p<0.10, ** p<0.05, *** p<0.01. Includes data from 512 Adj. CZs for 2005-2017.

Table A5: Housing supply: OLS baseline results

Dependent variable:	CZ Log House price growth$_{it}$				
Specification:	Average Effects		Heterogeneous Slopes		
	(1)	(2)	(3)	(4)	(5)
Δ Housing Expenditure	0.341***	0.351***	0.251***	0.237***	0.235***
	(0.03)	(0.03)	(0.02)	(0.04)	(0.04)
Δ Housing Exp. × Land unavail.			0.184***	0.200	0.194
			(0.06)	(0.13)	(0.13)
Δ Amenities index					0.769***
					(0.09)
Observations	10,773	10,773	13,338	13,338	13,338
10-yr Treas. Rates & Infl. Expect.	✓	✓	✓	✓	✓
CZ FE		✓		✓	✓

Heteroskedasticity-robust standard errors clustered at the CZ level in parentheses: * p<0.10, ** p<0.05, *** p<0.01. Includes data from 512 Adj. CZs for 2005-2017. "Housing expenditures" are the log of the sum of the wage bill in each education group multiplied by the group's calibrated housing expenditure share.

E Derivations

E.1 Static model estimating equation

In this section, I motivate the functional form for the reduced form analysis by deriving a static model of inter-city migration. This model shows why house price spillovers between cities are possible, and disciplines the functional form and empirical specification to test whether there are causal effects operating via migration networks. Moreover, this model introduces the key ideas underlying the discussion of identification concerns in the reduced-form estimation and motivates the choice of control variables to include in the estimation.

Worker preferences. All workers live in some city $i \in N$. They have Cobb-Douglas utility over tradable consumption goods with uniform prices across all locations, and local nontradable goods, including housing, with unit cost Q_{it}. The indirect utility for a worker ω in location i can therefore be written as:

$$U_{it}(\omega) = \ln A_{it} + \ln W_{it} - \alpha \ln Q_{it},$$

where α is the preference for nontradable goods. Here, W_{it} is the local wage and A_{it} captures residential amenities from living in i.

To model the heterogeneity among workers, I assume that in each period workers draw an idiosyncratic location amenity shock $z_i(\omega)$ for each location i, which is Type 1 extreme value (Gumbel) distributed with shape parameter θ. A smaller θ corresponds to less heterogeneity in idiosyncratic preferences among workers and therefore a higher sensitivity to differences in common factors between locations. The idiosyncratic location amenity is realized right after the moving decision.

In addition, whenever a worker moves between two locations i and k, she has to pay an additive and time-invariant moving cost τ^{ik} in utility units in the same period.

The worker ω chooses the destination location k by solving the following problem:

$$\max_{k \in N} \{U_{kt} - \tau^{ik} + z_k(\omega)\}.$$

Moreover, workers think of these moves as being once-and-for-all and don't consider the possibility of future moves or future realizations of wages or housing costs – an assumption that I relax later in the quantitative model.[102] The

[102] Alternatively, one could think of workers in the static model as being born in city i, deciding to move (or not) to some location k, and then dying at the end of the period to be replaced by a descendant who does not enter their parent's utility function.

probability of a worker choosing to move from i to k in period t therefore takes the standard logit form

$$\mu_t^{ik} = \frac{\exp(U_{kt} - \tau^{ik})^{\frac{1}{\theta}}}{\sum_j^N \exp(U_{jt} - \tau^{ij})^{\frac{1}{\theta}}}. \tag{31}$$

Production. City i has multiple industries ι. In each industry, firms produce tradable output $Y_{\iota,it}$ under perfect competition and constant returns to scale. For simplicity, I assume that that the production technology is Cobb-Douglas, with its components consisting of labor L_{it}, capital K_{it} and local productivity $X_{\iota,it}$:

$$\ln Y_{\iota,it} = \ln X_{\iota,it} + \lambda \ln L_{\iota,it} + (1-\lambda) \ln K_{\iota,it} \tag{32}$$

I follow Moretti (2011) and assume that there is an international capital market, which supplies capital elastically at marginal cost i_t. Using the first order conditions of the production function, this means that wages do not depend on the scale of production as

$$\ln W_{\iota,it} = \frac{1}{\lambda} \ln X_{\iota,it} - \frac{1-\lambda}{\lambda} \ln i_t + \frac{1-\lambda}{\lambda} \ln(1-\lambda)\lambda.$$

Workers are perfectly mobile across industries within a city, but choose locations based on the average wage in the city, which is given by the employment-weighted average across local industries:

$$W_{it} = \sum_\iota \left(\frac{L_{\iota,it}}{L_{it}} \right) W_{\iota,it}$$

I assume that local industry-level productivity can be decomposed into a national industry trend $\chi_{\iota t}$ and an idiosyncratic city-level component χ_{it}:

$$\ln X_{\iota,it} = \chi_{\iota t} + \chi_{it}$$

To a first-order approximation, log changes in wages relative to an equilibrium in the previous period can then be written as

$$\begin{aligned} \Delta \ln W_{it} &= \sum_\iota \tilde{\omega}_{\iota,i,t-1} \Delta \ln W_{\iota,it} \\ &= -\frac{1-\lambda}{\lambda} \Delta \ln i_t + \frac{1}{\lambda} \sum_\iota \tilde{\omega}_{\iota,i,t-1} (\Delta \chi_{\iota t} + \Delta \chi_{it}), \end{aligned}$$

where $\tilde{\omega}_{\iota,i,t-1} = \frac{L_{\iota,i,t-1} W_{\iota,i,t-1}}{L_{i,t-1} W_{i,t-1}}$ is the wage bill share of industry ι in the baseline period, which measures the exposure of the labor market to shocks in this industry. Note that this means that we can write local changes in log average wages as a "shift-share" function where the "shifts" are the labor demand shocks $\Delta \chi_{\iota t}$ and $\Delta \chi_{it}$ and the exposure shares consist of the local wage bill shares (Adao et al., 2019).

Housing supply. House prices are a function of city size, i.e.

$$P_{it} = k_i L_{it}^{\eta_i^H} \tag{33}$$

where η_i^H is a supply elasticity parameter that can vary across cities, and κ_i is a location-specific housing productivity parameter. Both of these parameters are assumed to be exogenously determined and time-invariant.

House price formation. House prices P_{it} are connected to housing costs Q_{it} through a valuation function $Q_{it} = f(P_{it}, \cdot)$, where the other arguments of the valuation function depend on home buyers' information set and inference.

In general, home buyers are pricing the current and expected future flow of rents from a home, i.e.

$$P_{it} = Q_{it} + \frac{E_t[Q_{t+1}]}{1+r_t} + \frac{E_t[Q_{t+2}]}{1+r_t} + \ldots$$

I assume that home buyers are "myopic" in the sense that they price houses as if rent growth is constant at some value $g_t = \bar{g} \;\; \forall t$ over time. As a result, house prices are

$$P_{it} = \left(\frac{1+r_t}{r_t - \bar{g}} \right) Q_{it}$$

and their changes are given by

$$\Delta \ln P_{i,t+1} = \Delta \ln Q_{i,t+1} + \alpha_t^{\mathrm{m}}$$

where $\alpha_t^{\mathrm{m}} = \ln\left(\frac{(1+r_{t+1})(r_t-\bar{g})}{(1+r_t)(r_{t+1}-\bar{g})}\right)$, and $\alpha_t^{\mathrm{m}} = 0$ if discount rates r_t are constant over time. This is an important benchmark case because it corresponds closely to the implicit assumption in most static spatial equilibrium models. For instance, in constructing time series of changes in housing costs, rent data are often imputed as a constant multiple of house prices (see, e.g. Diamond (2016); Ganong and Shoag (2017)), or that rents can be mapped directly into contemporaneous house prices using interest rates (e.g. (Piyapromdee et al., 2014)). Similarly, spatial models in economic geography and trade often employ models where prices for real estate clear the market each period but are not forward-looking, such that per-period housing cost and house prices are the same (e.g. Ahlfeldt et al. (2015); Caliendo et al. (2019); Sturm et al. (2020)).

In an extension of the baseline specification, I also relax this standard assumption and allow for the valuation term to depend on past local price changes to incorporate the possibility of extrapolation. Then, the valuation function can generically be written as

$$\Delta \ln Q_{it} = \delta_p \Delta \ln P_{it} + \delta_x \Delta \ln P_{i,t-1} + \alpha_t, \tag{34}$$

where in the baseline specification $\delta_x = 0$.

E.1.1 Static Spatial Equilibrium

Given previous period population distribution $\{L_{i,t-1}\}_{i=1}^N$, migration costs $\{\{\tau^{ik}\}_{i=1}^N\}_{k=1}^N$, valuation changes α_t, wages $\{\{W_{i,t-1}\}_{i=1}^N\}$, and previous period housing cost $\{Q_{i,t-1}, P_{i,t-1}\}_{i=1}^N$ (as well as house price growth $\{\Delta \ln P_{i,t-1}\}_{i=1}^N$ when allowing for extrapolation), and location amenities and productivities $\{A_{it}, X_{it}\}_{i=1}^N$, an equilibrium at time t in the static model is defined as a series of population leves $\{L_{it}\}_{i=1}^N$, wages $\{\{W_{it}\}_{i=1}^N\}$, and housing cost and house prices $\{Q_{it}, P_{it}\}_{i=1}^N$, such that

1. Workers choose their location optimally when given the chance to migrate, such that

$$\mu_t^{ik} = \frac{\exp(U_{kt} - \tau^{ik})^{\frac{1}{\theta}}}{\sum_j^N \exp(U_{jt} - \tau^{ij})^{\frac{1}{\theta}}}. \tag{35}$$

   where

$$U_{it}(\omega) = \ln A_{it} + \ln W_{it} - \alpha \ln Q_{it},$$

   and populations in each location evolve as

$$L_{it} = \sum_k \mu_t^{ki} L_{k,t-1}$$

2. Wages are determined by

$$\ln W_{it} = \frac{1}{\lambda} \ln X_{it} - \frac{1-\lambda}{\lambda} \ln i_t + \frac{1-\lambda}{\lambda} \ln(1-\lambda)\lambda$$

3. House prices are determined by the housing supply condition

$$P_{it} = k_i L_{it}^{\eta_i^H}$$

4. House prices are linked to rents through the valuation function

$$\Delta \ln Q_{it} = \delta_p \Delta \ln P_{it} + \delta_x \Delta \ln P_{i,t-1} + \alpha_t,$$

   where in the baseline specification $\delta_x = 0$.

In the benchmark case, we are solving for $N \times 4$ variables (wages, house prices, rent, and population for each city), and have N equations each from the four equilibrium conditions above, so the equilibrium is exactly identified.

E.1.2 Spillovers

We can use the migration choices to derive the expected spillovers between cities. First, note that we can take the log of the migration choice probabilities to obtain

$$\ln \mu_t^{ki} = \frac{1}{\theta}(U_{it} - \tau^{ki}) - \ln \left(\sum_j^N \exp(U_{jt} - \tau^{kj})^{\frac{1}{\theta}} \right).$$

Then, substituting the components of flow utility for U_{kt} and totally differentiating both sides, we obtain

$$\begin{aligned}\Delta \ln \mu_t^{ki} =& \theta^{-1}\Delta \ln A_i + \theta^{-1}\Delta \ln W_{it} - \theta^{-1}\alpha\Delta \ln Q_{it} \\ & - \sum_j \mu_t^{kj} \left(\theta^{-1}\Delta \ln A_j + \theta^{-1}\Delta \ln W_{jt} - \theta^{-1}\alpha \Delta \ln Q_{jt}\right).\end{aligned}$$

That is, any increase in average migration probabilities either comes from a change in the expected utility of the destination or a deterioration in the other available options.

This means that we can write the change in local population as

$$\begin{aligned}\Delta \ln L_{it} =& \Delta \ln \left(\sum_k \mu_t^{ki} L_{k,t-1} \right) \\ =& \sum_k \phi_{t-1}^{i\leftarrow k} \left(\Delta \ln \mu_t^{ki} + \Delta \ln L_{k,t-1}\right)\end{aligned}$$

where $\phi_{t-1}^{i\leftarrow k} = \mu_{t-1}^{ki} L_{k,t-2} / \sum_j \mu_{t-1}^{ji} L_{j,t-2}$ is the share of city i's population coming from city k in the previous period. These terms represent city i's exposure to changes in other cities.

Now, we can substitute for the log change in migration shares from above to write

$$\begin{aligned}\Delta \ln L_{it} = & \underbrace{\sum_k \phi_{t-1}^{i\leftarrow k} \Delta \ln L_{k,t-1}}_{\text{Migration access: } \Delta \mathcal{M}_{i,t-1}} + \theta^{-1} \cdot \Delta \ln A_{it} + \theta^{-1} \cdot \Delta \ln W_{it} - \alpha\theta^{-1} \Delta \ln Q_{it} \\ & + \theta^{-1} \left(\alpha \cdot \underbrace{\Delta \widetilde{\mathcal{Q}}_{it}}_{\text{Network housing costs}} - \underbrace{\Delta \widetilde{\mathcal{A}}_{it}}_{\text{Network amenities}} - \underbrace{\Delta \widetilde{\mathcal{W}}_{it}}_{\text{Network wages}} \right)\end{aligned} \tag{36}$$

where the notation $\Delta \widetilde{\mathcal{X}}_{it} = \sum_k \phi_{t-1}^{i\leftarrow k} \sum_j \mu_{t-1}^{kj} \Delta \ln X_{jt}$ denotes a migration-weighted sum of the log change in characteristic X over city i's entire migration network.

Substituting equation 34 into the definition of the network housing cost term $\Delta \widetilde{\mathcal{Q}}_{it}$ we can decompose the network term as follows:

$$\begin{aligned}\Delta \widetilde{\mathcal{Q}}_{it} =& \sum_k \phi_{t-1}^{i\leftarrow k} \sum_j \mu_{t-1}^{kj} \Delta \ln Q_{jt} \\ =& \alpha_t + \zeta_{t-1}^i (\delta_p \Delta \ln P_{it} + \delta_x \Delta \ln P_{i,t-1}) \\ & + \delta_p \cdot \underbrace{\left(\sum_{k\in N} \phi_{t-1}^{i\leftarrow k} \sum_{j:j\neq i} \mu_{t-1}^{kj} \Delta \ln P_{jt} \right)}_{\text{Network House Price Effect:} \Delta \mathcal{P}_{it}^{NW}} + \delta_x \cdot \underbrace{\left(\sum_{k\in N} \phi_{t-1}^{i\leftarrow k} \sum_{j:j\neq i} \mu_{t-1}^{kj} \Delta \ln P_{j,t-1} \right)}_{\text{Network Extrapolation Effect: } \Delta \mathcal{P}_{i,t-1}^{NW}} \\ =& \alpha_t + \zeta_{t-1}^i \delta_p \Delta \ln P_{it} + \zeta_{t-1}^i \delta_x \Delta \ln E[P_{i,t+1}] + \delta_p \Delta \mathcal{P}_{i,t}^{NW} + \delta_x \Delta \mathcal{P}_{i,t-1}^{NW}\end{aligned}$$

where $\zeta_{t-1}^i = \sum_{k\in N} \phi_{t-1}^{i\leftarrow k} \mu_{t-1}^{ki}$ captures indirect migration network exposure of the focal city to its own price changes.

We can substitute both the network housing cost and regular housing cost expressions into the population change equation 36:

$$\begin{aligned}\Delta\ln L_{it} =& \Delta\mathcal{M}_{i,t-1} + \theta^{-1}\cdot\Delta\ln A_{it} + \theta^{-1}\cdot\Delta\ln W_{it} - \alpha\theta^{-1}\Delta\ln Q_{it}\\ &+ \theta^{-1}\alpha\Delta\widetilde{\mathcal{Q}}_{i,t} - \theta^{-1}\Delta\widetilde{\mathcal{A}}_{it} - \theta^{-1}\Delta\widetilde{\mathcal{W}}_{it}\\ =& \Delta\mathcal{M}_{i,t-1} + \theta^{-1}\left(\Delta\ln A_{it} + \Delta\ln W_{it} - \Delta\widetilde{\mathcal{A}}_{it} - \Delta\widetilde{\mathcal{W}}_{it}\right) - \alpha\theta^{-1}\left(\delta_p\ln P_{it} + \delta_x\Delta\ln E[P_{i,t+1}] + \alpha_t\right)\\ &+ \theta^{-1}\alpha\left(\alpha_t + \zeta^i_{t-1}\delta_p\Delta\ln P_{it} + \zeta^i_{t-1}\delta_x\Delta\ln P_{i,t-1} + \delta_p\Delta\mathcal{P}^{NW}_{i,t} + \delta_x\Delta\mathcal{P}^{NW}_{i,t-1}\right)\\ =& \Delta\mathcal{M}_{i,t-1} - \alpha\theta^{-1}(1-\zeta^i_{t-1})\left(\delta_p\Delta\ln P_{i,t} + \delta_x\Delta\ln P_{i,t-1}\right) + \alpha\theta^{-1}\delta_p\mathcal{P}^{NW}_{it}\\ &+ \alpha\theta^{-1}\delta_x\mathcal{P}^{NW}_{i,t-1} + \xi^L_{it},\end{aligned}$$

where $\xi^L_{it} = \theta^{-1}\left(\Delta\ln A_{it} + \Delta\ln W_{it} - \Delta\widetilde{\mathcal{A}}_{it} - \Delta\widetilde{\mathcal{W}}_{it}\right)$. In words, equilibrium population changes in city i depend negatively on changes in the cost of living in i, and positively on the cost of living in other cities in i's migration network, as well as any changes in the relative attractiveness of city i due to amenity or productivity changes.

Next, we can substitute the inverse housing supply curve into this population change equation, allowing for heterogeneous inverse housing supply elasticity η^H_i. Then, isolating $\Delta\ln L_{it}$, we can get an expression of city i population change in terms of its own and network characteristics:

$$\Delta\ln L_{it} = \eta^M_i\Delta\mathcal{M}_{i,t-1} + \eta^x_i\Delta\ln P_{i,t-1} + \eta^{nw}_i\Delta\mathcal{P}^{NW}_{it} + \eta^{nwx}_i\Delta\mathcal{P}^{NW}_{i,t-1} + \tilde{\xi}^L_{it}$$

where

$$\begin{aligned}\eta^M_i =& (1 + \eta^H_i\alpha\theta^{-1}(1-\zeta^i_{t-1})\delta_p)^{-1}\\ \eta^x_i =& -\eta^M_i\alpha\theta^{-1}(1-\zeta^i_{t-1})\delta_x\\ \eta^{nw}_i =& \eta^M_i\alpha\theta^{-1}\delta_p\\ \eta^{nwx}_i =& \eta^M_i\alpha\theta^{-1}\delta_x\\ \tilde{\xi}^L_{it} =& \eta^M_i\theta^{-1}\left(\Delta\ln A_{it} + \Delta\ln W_{it} - \Delta\widetilde{\mathcal{A}}_{it} - \Delta\widetilde{\mathcal{W}}_{it}\right)\end{aligned}$$

Note that the coefficients here might differ across cities for two reasons: On the one hand, changes in the network house prices will have a greater effect on population in cities that have a higher gross migration activity to begin with, i.e. a smaller ζ^i_{t-1}. On the other hand, the population change that results in equilibrium from any change in relative city attractiveness is smaller the greater the inverse housing supply elasticity η^H_i. That is, cities where the housing supply is inelastic see smaller migration responses for given network changes.

Now, we can simply plug this expression into $\Delta\ln P_{it} = \eta^H_i\Delta\ln L_{it}$, to find the equilibrium house price growth:

$$\Delta\ln P_{it} = \eta^H_i\left(\eta^M_i\Delta\mathcal{M}_{i,t-1} + \eta^x_i\Delta\ln P_{i,t+x} + \eta^{nw}_i\Delta\mathcal{P}^{NW}_{it} + \eta^{nwx}_i\Delta\mathcal{P}^{NW}_{i,t-1} + \tilde{\xi}^L_{it}\right)$$

In the baseline estimation the effect of heterogeneity in inverse housing supply elasticity across cities is assumed to be exogenous with regard to other explanatory variables such that I can gather its effect into the residual, but I will also present results that explicitly allows coefficients to be variable across cities.

Thus, the main reduced form estimating equation becomes

$$\Delta\ln P_{it} = \tilde{\eta}^M\Delta\mathcal{M}_{i,t-1} + \tilde{\eta}^x\Delta\ln P_{i,t-1} + \tilde{\eta}^{nw}\Delta\mathcal{P}^{NW}_{it} + \tilde{\eta}^{nwx}\Delta\mathcal{P}^{NW}_{i,t-1} + \tilde{\alpha}^P_t + \xi^P_{it},$$

where

$$\begin{aligned}\tilde{\eta}^M =& E[\eta^H_i\eta^M_i]\\ \tilde{\eta}^x =& E[\eta^H_i\eta^x_i]\\ \tilde{\eta}^{nw} =& E[\eta^H_i\eta^{nw}_i]\\ \tilde{\eta}^{nwx}_i =& E[\eta^H_i\eta^{nwx}_i]\\ \xi^P_{it} =& \eta^H_i\eta^M_i\theta^{-1}\left(\Delta\ln A_{it} + \Delta\ln W_{it} - \Delta\widetilde{\mathcal{A}}_{it} - \Delta\widetilde{\mathcal{W}}_{it}\right) + (\eta^H_i\eta^M_i - E[\eta^H_i\eta^M_i])\Delta\mathcal{M}_{i,t-1}\\ &+ (\eta^H_i\eta^x_i - E[\eta^H_i\eta^x_i])\Delta\ln P_{i,t+x} + (\eta^H_i\eta^{nw}_i - E[\eta^H_i\eta^{nw}_i])\Delta\mathcal{P}^{NW}_{it} + (\eta^H_i\eta^{nwx}_i - E[\eta^H_i\eta^{nwx}_i])\Delta\mathcal{P}^{NW}_{i,t-1}.\end{aligned}$$

In the baseline model without extrapolation, $\tilde{\eta}^{x} = \tilde{\eta}^{nwx} = 0$.

In order to capture potentially confounding trends included in the error term ξ_{it}^{P}, I can explicitly control for a number of covariates Γ_{it}, including $\Delta\mathcal{M}_{i,t-1}$, as discussed in Section 3.5. Then, the baseline estimating equation is

$$\Delta \ln P_{it} = \alpha_i + \alpha_t + \tilde{\eta}^{nw}\Delta\mathcal{P}_{it}^{NW} + \beta'\Gamma_{it} + \tilde{\xi}_{it}^{P},$$

where $\tilde{\xi}_{it}^{P} = \xi_{it}^{P} - \alpha_i - \alpha_t - \beta'\Gamma_{it}$ captures price residuals after controlling for the confounders.

Intuitively, we solved for the increase in local house prices as a result of population changes that is consistent with location choices. That is, after observing the change in house prices in city i in response to period t migration, the marginal migrant is indifferent between city i and their origin city. Note that migrants in this model are not "surprised" by the house prices changes in their destination city in response to migration, but rather take the response of house prices in their destination city into account when deciding to move.

Identification issues arise because residual changes in house prices $\tilde{\xi}_{it}^{P}$ may be driven by residual changes in local amenities or labor productivity relative to the migration network, as well as other idiosyncratic valuation shocks, even after controlling for the covariates discussed in Section 3.5.

E.2 Network IV exclusion restriction derivation

Shock-level reformulation of network IV. Following Borusyak et al. (2020), I can formulate the exclusion restriction at the level of industry shocks. For simplicity, I focus on just one instrument $z_{it} = NWP_{it}^{B}$, but the analysis generalizes to the case of multiple instruments. The Frisch-Waugh Lovell Theorem implies that the spillover coefficient estimate in a panel of length T with N cities can be written as the second-stage coefficient of a residualized IV regression of the form

$$\hat{\tilde{\eta}}^{nw} = \frac{\sum_{t=1}^{T}\sum_{i=1}^{N} z_{it} y_{it}^{\perp}}{\sum_{t=1}^{T}\sum_{i=1}^{N} z_{it} x_{it}^{\perp}}$$

where $y_{it} = \Delta \ln P_{it}$ for the main house price spillover estimation, and $x_{it} = \Delta\mathcal{P}_{it}^{NW}$. The notation $\nu_{it}^{\perp}$ denotes the residual from a projection of ν_{it} on the vector of control variables Γ_{it} detailed in Section 3.5. Combining the definition of the network instruments in Equation 6 and the shift-share shock definition, we can write

$$\begin{aligned} z_{it} &= \sum_{j:j\neq i} \psi_{\text{`90-`95}}^{ij} \sum_{\iota=1}^{N_{ind}} \tilde{\omega}_{\iota,j,\text{`90}} \Delta \ln W_{\iota,-j,t}^{US} \\ &= \sum_{\iota=1}^{N_{ind}} s_{i\iota} g_{\iota,t} \end{aligned}$$

where $g_{\iota t} \approx \Delta \ln W_{\iota,-j,t}^{US}\ \forall j$, ignoring the fact that the industry shock is constructed as a leave-one-out variable and therefore varies slightly between cities in a finite sample.[103] Here, I have combined the industry exposure and migration network structure into a weight $s_{i\iota} = \sum_{j:j\neq i} \psi_{\text{`90-`95}}^{ij} \tilde{\omega}_{\iota,j,\text{`90}}$ that summarizes the migration network exposure of city i to industry ι. We can then write

$$\hat{\tilde{\eta}}^{nw} = \frac{\sum_{t=1}^{T}\sum_{\iota=1}^{N_{ind}} s_{\iota} g_{\iota t} \bar{y}_{\iota t}^{\perp}}{\sum_{t=1}^{T}\sum_{\iota=1}^{N_{ind}} s_{\iota} g_{\iota t} \bar{x}_{\iota t}^{\perp}},$$

where the notation $\bar{\nu}_t = \frac{\frac{1}{N}\sum_{i=1}^{N} s_{i\iota}\nu_{it}}{\frac{1}{N}\sum_{i=1}^{N} s_{i\iota}}$ represents a weighted average across cities, with weights given by cities' relative exposure to industry ι through their migration network.[104] Moreover, $s_{\iota} = \frac{1}{N}\sum_{i=1}^{N} s_{i\iota}$ is the average city migration exposure to industry ι growth. This expression shows that we could estimate the same IV spillover coefficient in an industry-level regression of industry exposure weighted city house price growth on industry exposure weighted network house price changes, using national industry growth trends as the instrument.

[103] This approximation is asymptotically correct as the number of cities gets large, and simplifies the notation and intuition substantially here. The main insight is robust to using the leave-one-out version.

[104] To see how I arrived at this expression, note that the numerator can be written as $\sum_{t=1}^{T}\sum_{i=1}^{N} z_{it} y_{it}^{\perp} = \sum_{t=1}^{T}\sum_{i=1}^{N}\sum_{\iota=1}^{N_{ind}} s_{i\iota} g_{\iota,t} y_{it}^{\perp} = \sum_{t=1}^{T}\sum_{\iota=1}^{N_{ind}} s_{\iota} g_{\iota,t} \left(\frac{\sum_{i=1}^{N} s_{i\iota} y_{it}^{\perp}}{\sum_{i=1}^{N} s_{i\iota}}\right) = \sum_{t=1}^{T}\sum_{\iota=1}^{N_{ind}} s_{\iota} g_{\iota t} \bar{y}_{\iota t}^{\perp}$. Rewriting the denominator in a similar way gives the expression shown.

This rewriting of the estimator at the industry level provides a perspective on the identifying variation underlying the network IV estimate: We can think of the spillover effect estimation as identifying the spillover effect from the degree to which the covariance in industry wage growth shocks with house price growth in the cities most exposed to the industry (the numerator) is higher in the cities that are "treated" in the form of having house price changes in their migration network that vary with industry wage shocks (the denominator). For instance, if cities with migration connections to technology hubs (e.g. Boise, ID) see higher house price growth precisely when tech hub house prices rise (e.g. in San Francisco) as a result of national trends in knowledge industry wages, then this variation allows us to infer a causal positive spillover effect.

Network IV exclusion restriction. This rewriting of the network IV estimator in the form of industry-level shocks then allows me to formulate the exclusion restriction of the network approach as follows:

Proposition 1 *If the network instrument NWP_{it}^{B} (or NWP_{it}^{Bx}) is relevant and mild regularity conditions hold (the variance matrix of control variables has full rank, and the covariance matrices of instruments and residuals with controls are bounded and exist), then the IV estimate of the spillover effect $\hat{\tilde{\eta}}^{nw}$ is consistent if and only if*

$$\sum_{t=1}^{T}\sum_{\iota=1}^{N_{ind}} s_\iota g_{\iota t}\bar{\xi}_{\iota t}^{P,\perp} \to^p 0.$$

The proof is analogous to that for Proposition 2 in Borusyak et al. (2020). Here, $\bar{\xi}_{\iota,t}^{P,\perp}$ is the error in the house price growth regression in Equation 4, residualized with regard to the control variables Γ_{it}, and averaged over cities, weighting them by their migration network exposure to industry shocks.

This condition shows that, for the network IV estimate to be consistent, industry wage shocks cannot be systematically higher for those industries that have a systematically larger migration network impact on cities that are experiencing large unobserved house price shocks, conditional on control variables.

As Borusyak et al. (2020) argue, this identification allows for a city's migration network to be endogenously determined – it only requires the national industry trends over time to be exogenous in the sense defined in Proposition 1. This would be invalidated, if, for example, cities that experience more migration flows from cities that specialize in the booming tech industry are also systematically experiencing greater idiosyncratic house price movements in a way that is not captured by their own exposure to knowledge industries or any other included control variables.

This approach also provides insights regarding the concern that the network house price instruments are correlated with focal city i industry shocks if industry structure is correlated across cities that share migration links. Note that this concern is supported by the significant coefficient on industry structure in the migration cost determinants analyzed in Section 2.2: Table 2 showed that migration costs appear to be lower among cities with similar industry structures, making them more likely to have strong migration links.

I can explicitly express this concern in the industry shock formulation of the exclusion restriction: if the direct shift-share shock B_{it} is not included in the control variables, and industry wage shocks affect city i house prices, then the house price residual $\xi_{it}^{P,\perp}$ would be a function of $B_{it} = \sum_\iota \tilde{\omega}_{\iota,i,`90}^{\perp} g_{\iota,t}$. Using the definition of $\bar{\xi}_{\iota,t}^{P,\perp}$, we can write

$$\bar{\xi}_{\iota,t}^{P,\perp} = \frac{\sum_{i=1}^{N} s_{i\iota}\xi_{it}^{P,\perp}}{\sum_{i=1}^{N} s_{i\iota}} = \frac{\sum_{i=1}^{N}\sum_{j:j\neq i}^{N} \psi_{`90\text{-}`95}^{ij}\tilde{\omega}_{\iota j,`90}\xi_{it}^{P,\perp}}{\sum_{i=1}^{N} s_{i\iota}} = \frac{\sum_{i=1}^{N}\sum_{j:j\neq i}^{N} \psi_{`90\text{-}`95}^{ij}\tilde{\omega}_{\iota j,`90}(\sum_\iota \tilde{\omega}_{\iota i,`90}^{\perp} g_{\iota,t} + \epsilon_{it})}{\sum_{i=1}^{N} s_{i\iota}}.$$

Note that this expression contains $\sum_{i=1}^{N}\sum_{j:j\neq i}^{N} \psi_{`90\text{-}`95}^{ij}\tilde{\omega}_{\iota j,`90}\tilde{\omega}_{\iota i,`90}^{\perp}$, which is the average migration-weighted covariance in (demeaned) industry employment shares of a city with other cities. If a city is more likely to receive migrants from other cities that are likely to have large industry employment shares in the same industries, then this covariance will be positive, and the exclusion restriction will not hold.

Dynamic effect exclusion restriction. I also estimate IV forecasting regressions that correspond to local projections with external instruments (Jordà, 2005; Stock and Watson, 2018) of the form

$$\Delta \ln P_{i,t-1+h} = \alpha_i + \alpha_t + \tilde{\eta}_h^{nw}\Delta\mathcal{P}_{it}^{NW} + \beta'\Gamma_{it} + \xi_{i,t-1+h}^{P}, \tag{37}$$

where the vector Γ_{it} includes the same additional control variables as the static model (see Section 3.5), and I am again instrumenting for the time t network price growth shock using network labor demand shocks. The coefficients $\tilde{\eta}_h^{nw}$ now represent the impulse response in period $t-1+h$ of the shock. That is, the contemporaneous impacts correspond to $h = 1$, the impact on the dependent variable in the year after the shock is $h = 2$, and so on.

As Stock and Watson (2018) note, to interpret $\tilde{\eta}_h^{nw}$ as causal dynamic effects we need the period t network instruments to not just be uncorrelated with contemporaneous price growth shocks $\xi_{i,t}^P$, but also all leads and lags of the shock. That is, the exclusion restriction becomes

$$\sum_{t=1}^{T}\sum_{\iota=1}^{N_{ind}} s_\iota g_{\iota t}\bar{\xi}_{\iota,t-1+h}^{P,\perp} \to^p 0. \ \forall h.$$

This condition requires that industry wage shocks are not predictable given the exposure weighted pattern of city house price growth residuals in *any* period. This is a stronger condition than for the contemporaneous effect. It would be violated, for instance, if house prices rise in anticipation of a technology boom in other cities that share migration links.

However, this restriction again only needs to hold conditional on the included control variables. I control, for example, for contemporaneous wage changes and productivity shocks in the focal city and its migration network, as well as regional trends. This means that, to the degree that the economic dynamics anticipated by past house price run-ups in the focal city are captured by contemporaneous control variables, industry shocks will be exogenous *conditional* on including these covariates.

E.3 Quantitative model derivations

E.3.1 Labor demand equation derivation

Production technology. Firms in each industry ι are indifferent between workers of any group s and produce output $Y_{i\iota t}$ that consists of tradable goods differentiated by industry. The production technology uses local labor as the only input and has constant returns to scale:

$$Y_{i\iota t} = X_{i\iota t} L_{i\iota t},$$

where $L_{i\iota t} = \sum_s^{\mathcal{S}} L_{i\iota ts}$ is the sum of local industry ι employment across workers in groups $s \in \mathcal{S}$. Here, $X_{i\iota t}$ denotes the local productivity of firms producing goods in industry ι. Moreover, I assume that there is perfect competition in input markets such that workers earn their marginal product.

Tradable goods demand. I model the local demand for differentiated tradable goods following Armington (1969).[105] I assume that, within the tradables category, the consumption utility in location i over which different groups s have Cobb-Douglas preferences is an aggregator of local industry good consumption of the form

$$C_{it} = \prod_{\iota=1}^{\mathcal{I}} C_{i\iota t}^{\gamma_\iota}, \tag{38}$$

where $C_{i\iota t}$ – the industry ι goods consumption in city i – is in turn an aggregator of the utility obtained from the consumption of differentiated goods from other locations j given by

$$C_{i\iota t} = \left(\sum_{j=1}^{N}(c_{ij\iota})^{\frac{\sigma-1}{\sigma}}\right)^{\frac{\sigma}{\sigma-1}},$$

where $\sigma \in (1,\infty)$ and $c_{ij\iota}$ is the consumption of industry ι goods from city j in city i. Trade is assumed to be costless and product markets are competitive, so goods from a particular location cost the same in all other locations. Therefore, we can write

$$P_{\iota t}^{1-\sigma} = \sum_{i=1}^{N}(p_{i\iota t})^{1-\sigma}$$

as the composite price index for goods from industry ι.

Based on the consumer preferences, the share of all tradable goods consumption spending going to city i goods in industry ι takes the form (Armington, 1969)

$$x_{i\iota t} = \left(\frac{p_{i\iota t}}{P_{\iota t}}\right)^{1-\sigma}. \tag{39}$$

[105] This demand structure is common in the trade and economic geography literature – see, e.g. Arkolakis et al. (2012), or Adao et al. (2019) for a similar recent application.

Product and labor market equilibrium. Competitive output markets imply that the output price equals marginal costs in each local industry, which are given by productivity adjusted wages, as labor is the only production input:

$$p_{i\iota t} = \frac{W_{i\iota t}}{X_{i\iota t}} \tag{40}$$

As firms earn no profit, total revenues in each local goods category have to equal the total wage bill $W_{i\iota t}L_{i\iota t}$ in that sector. At the same time, revenues in industry ι and location i correspond to a share $x_{i\iota t}\gamma_\iota$ of global demand, which I normalize to one. Substituting for the demand share from Equation 39 and for local prices from the marginal cost Equation 40, this implies that the local industry labor demand equation can be written as

$$L_{i\iota t} = W_{i\iota t}^{-\sigma} D_{i\iota t},$$

where $D_{i\iota t} = (X_{i\iota t}P_{\iota t})^{\sigma-1}\gamma_\iota$ summarizes local industry demand and productivity shocks.

Rearranging, taking logs, and first-differencing around the equilibrium, changes in log inverse labor demand can then be written as

$$\Delta \ln W_{i\iota t} = -\frac{1}{\sigma}\Delta \ln L_{i\iota t} + \frac{1}{\sigma}\Delta \ln D_{i\iota t},$$

To motivate a shift-share instrument approach, note that we can assume that local demand shocks can be decomposed into a national industry component and an idiosyncratic local shock as $D_{i\iota t} = D_{\iota t}\tilde{D}_{i\iota t}$. A shift-share approach based on the effect of national industry trends on local wage changes in this model then follows directly from proxying for national industry shocks using observed national wage changes.

Agglomeration. We can also take into account the possibility of agglomeration effects from greater employment, which have been previously documented in the literature (Diamond, 2016). Assume that

$$X_{i\iota t} = \widetilde{X}_{i\iota t} L_{i\iota t}^{\alpha},$$

so productivity in each industry changes with the local employment level, and $\widetilde{X}_{i\iota t}$ is the residual local industry productivity. Then, local industry wage changes become

$$\Delta \ln W_{i\iota t} = \eta^{LD}\Delta \ln L_{i\iota t} + \frac{1}{\sigma}\Delta \ln \tilde{D}_{i\iota t},$$

where $\widetilde{D}_{i\iota t} = \left(\widetilde{X}_{i\iota t}P_{\iota t}\right)^{\sigma-1}\gamma_\iota$, and $\eta^{LD} = \left(\frac{\alpha(\sigma-1)-1}{\sigma}\right)$. This is the labor demand equation shown in Equation 17 that is taken to the data.

E.3.2 House price formation: discussion of assumptions

House price formation. House prices P_{it} are connected to housing costs Q_{it} through a valuation function which depends on home buyers' information set. In general, home buyers are pricing the current and expected future flow of rents from a home, i.e.

$$P_{it} = Q_{it} + \frac{E_t[Q_{t+1}]}{1+r_t} + \frac{E_t[Q_{t+2}]}{1+r_t} + \dots$$

Moreover, from a homeowner perspective, housing costs in future equilibria will evolve from current housing costs based on a predictable growth rate g_{it} and an unforecastable shock ϵ_{it}^g:

$$Q_{i,t+1} = Q_{it}(1+g_{it})(1+\epsilon_{it}^g).$$

In the context of the model, g_{it} may capture, for instance, trends in the productivity and amenity changes driving location choices, while ϵ_{it}^g reflects randomness introduced by idiosyncratic amenity shock realizations, or shocks to the productivity and amenity terms. As long as home buyers know g_{it} in advance, house price changes should not be forecastable once we control for growth in the sources of fundamental location demand. However, as Glaeser and Nathanson (2017) note, rational pricing models can impose a significant cognitive load on homebuyers who are trying to infer the average willingness to pay Q_{it} or the growth rate process g_{it} from noisy data about past prices.

Rational home buyers. If home buyers are fully rational, they would have full information about the current forecastable expected growth rate of location fundamentals g_i and the present value and history of housing cost $\{Q_{it}, Q_{i,t-1}, Q_{i,t-2}, \dots\}$. House prices are then given by

$$P_{it} = \left(\frac{1+r_t}{r_t - g_{it}}\right) Q_{it}.$$

As a result, their valuation of houses only changes between periods as the result of local shocks ϵ^g_{it} or changes in discount rates r_t or the expected growth rate of fundamentals. Under rational pricing, log changes in house prices can therefore be written as

$$\Delta \ln P_{i,t+1} = \Delta \ln Q_{i,t+1} + \alpha_{it},$$

where $\alpha_{it} = \ln\left(\frac{(1+r_{t+1})(r_t-g_{it})}{(1+r_t)(r_{t+1}-g_{i,t+1})}\right)$.[106]

Myopic home buyers. If instead home buyers are "myopic" in the sense that they price houses as if rent growth is constant at some value $g_t = \bar{g}\ \ \forall t$ over time, log changes in house prices are given by

$$\Delta \ln P_{i,t+1} = \Delta \ln Q_{i,t+1} + \alpha^{\text{my}}_t$$

where $\alpha^{\text{my}}_t = \ln\left(\frac{(1+r_{t+1})(r_t-\bar{g})}{(1+r_t)(r_{t+1}-\bar{g})}\right)$. Note, however, that the myopic pricing case would have the same implication for the relationship between house price growth and rent growth as the rational pricing model, if forecastable housing demand changes can vary over time, but not across cites, i.e. $g_{it} = g_t\ \ \forall i$.

E.3.3 Housing supply function

Housing developers are assumed to be price-takers who sell housing consumption units at marginal production cost. Individual developers are marginal with regard to competition, but their output as a whole can be written in the form of a representative firm. Housing units are supplied using a Cobb-Douglas production technology that combines perfectly mobile construction capital K^c with units of local land A^c and local construction productivity κ_i:

$$H_{it} = \kappa_i (K^c_{it})^{\alpha_K} (A^c)^{\alpha_A}$$

Housing construction has constant returns to scale in land and capital, i.e. $\alpha_A + \alpha_K = 1$. The developers rent capital equipment at competitive national rates r^K_t, and acquire land at local land cost LC_{it} per unit.

The cost of land is assumed to increase with the size of the overall housing stock with a location-specific elasticity:

$$LC_{it} = H_{it}^{\alpha^i_L}.$$

The developer, however, does not take these general equilibrium effects on productivity into account in her optimization problem - she treats land costs as given.

Developers choose the scale of the housing stock to maximize profits from the sale of housing units, taking productivity as given. That is, developers solve

$$\max_{H_{it}} \{H_{it} P_{it} - LC_{it} A^c_{it} - r^K_t K^c_{it}\}.$$

[106] In detail, the log change in house prices from the perspective of a rational home buyer can be rewritten as

$$\begin{aligned}
\Delta \ln P_{i,t+1} &= \ln\left(\frac{Q_{i,t+1} + \frac{E_{t+1}[Q_{t+2}]}{1+r_{t+1}} + \dots}{Q_{it} + \frac{E_t[Q_{t+1}]}{1+r_t} + \frac{E_t[Q_{t+2}]}{(1+r_t)^2} + \dots}\right) \\
&= \ln\left(\frac{Q_{it}(1+g_{it})(1+\epsilon^g_{it}) + \frac{Q_{it}(1+g_{it})(1+\epsilon^g_{it})(1+g_{i,t+1})}{1+r_{t+1}} + \dots}{Q_{it} + \frac{Q_{it}(1+g_{it})}{1+r_t} + \frac{Q_{it}(1+g_{it})^2}{(1+r_t)^2} \dots}\right) \\
&= \Delta \ln Q_{i,t+1} + \ln\left(\frac{(1+r_{t+1})(r_t - g_{it})}{(1+r_t)(r_{t+1} - g_{i,t+1})}\right).
\end{aligned}$$

subject to the zero profit condition $H_{it}P_{it} - LC_{it}A^c_{it} - r^K_t K^c_{it} = 0$. Solving the corresponding cost minimization problem with regard to the input amounts, I find that the developer's cost function is

$$C_{it}(H_{it}) = LC_{it}^{\frac{\alpha_A}{\alpha_A+\alpha_K}} H_{it}^{\frac{1}{\alpha_A+\alpha_K}} \kappa_i^{\frac{-1}{\alpha_A+\alpha_K}} (\alpha_A+\alpha_K)\left(r_t^{\alpha_K}\alpha_A^{-\alpha_A}\alpha_K^{-\alpha_K}\right)^{\frac{1}{\alpha_A+\alpha_K}}$$

Note that the developer does not behave as a social planner and therefore does not take into account the indirect impact of new construction on costs through higher land costs.

The marginal cost function from the perspective of the developer is therefore

$$MC_{it} = \frac{\partial C_{it}(H_{it})}{\partial H_{it}} = LC_{it}^{\frac{\alpha_A}{\alpha_A+\alpha_K}} H_{it}^{\frac{1-\alpha_A-\alpha_K}{\alpha_A+\alpha_K}} \kappa_i^{\frac{-1}{\alpha_A+\alpha_K}} \left(r_t^{\alpha_K}\alpha_A^{-\alpha_A}\alpha_K^{-\alpha_K}\right)^{\frac{1}{\alpha_A+\alpha_K}}$$

However, for the city overall, land costs are endogenous, so the housing supply function for the city as a whole corresponds to the representative developer's marginal cost but taking into account the endogeneity of LC_{it}. Substituting for LC_{it}, and imposing $P_{it} = MC_{it}$, I find

$$P_{it} = \phi_i \phi_t H_{it}^{\phi_i^H}, \tag{41}$$

where

$$\phi_H^i = \frac{\alpha_L^i \alpha_A + 1 - \alpha_K - \alpha_A}{\alpha_K + \alpha_A}$$
$$\phi_i = \kappa_i^{\frac{-1}{\alpha_K+\alpha_A}}$$
$$\phi_t = \left((r_t^K)^{\alpha_K}\alpha_A^{-\alpha_A}\alpha_K^{-\alpha_K}\right)^{\frac{1}{\alpha_K+\alpha_A}}$$

Equation (41) describes the inverse housing supply curve. The parameter ϕ_H^i describes the inverse housing supply elasticity, and can vary across cities depending on local land constraints.

E.4 Moments and exclusion restrictions for the structural parameter estimation

This section complements Section 6 by providing additional details on the moments used and the exclusion restrictions that identify the structural parameters which are estimated using IV approaches.

E.4.1 Location choice parameter θ_s

The location choice sensitivity to differences in observed city characteristics is estimated from the equation

$$\Delta \mathbb{M}_{st}^{ik} = \theta_s^{-1}\Delta \ln\left(\frac{R_{kst}^w}{R_{ist}^w}\right) + \theta_s^{-1}\beta_s(\Delta ColShare_{kt} - \Delta ColShare_{it}) + \Delta\tilde{\xi}_t^{kis,Col}.$$

using shift-share instruments. The exclusion restriction for the shift-share instruments and the land constraints as a measure of housing supply slopes is that

$$Z_{ikst}^{lc} \perp \Delta\tilde{\xi}_t^{kis,Col} \;\; \forall [s \times (i,k)].$$

where

$$Z_{ikst}^{lc} = \begin{pmatrix} B_{it,\text{`}00,2d}, B_{it,\text{`}00,3d}, B_{it,\text{`}00}^{nc}, B_{it,\text{`}00}^{col}, B_{it,\text{`}00,2d}x_i^{land}, B_{it,\text{`}00,3d}x_i^{land} \\ B_{kt,\text{`}00,2d}, B_{kt,\text{`}00,3d}, B_{kt,\text{`}00}^{nc}, B_{kt,\text{`}00}^{col}, B_{kt,\text{`}00,2d}x_i^{land}, B_{kt,\text{`}00,3d}x_k^{land} \end{pmatrix}.$$

Under this assumption, we can identify θ_s^{-1} for college and non-college workers from the moments

$$E[Z_{ikst}^{lc} \cdot \Delta\tilde{\xi}_t^{kis,Col}] = 0.$$

E.4.2 Housing supply elasticities

The key estimating equation for the housing supply elasticity parameters is

$$\Delta \ln P_{it} = \eta^r \Delta \ln r_t^{mtg} + \eta^e \Delta \ln ei_t + \eta_i^H + (\psi_H + \psi_H^l x_i^{land}) \Delta \ln HD_{it} + \Delta \epsilon_{it}^P.$$

We can again use Bartik wage shift-shares as instruments to isolate exogenous variation in housing demand. The identifying assumption for the parameters of the housing supply function then becomes

$$\{B_{it,'00,3d}, B_{it,'00,3d} x_i^{land}\} \perp \Delta \epsilon_{it}^P$$

If this condition holds, we can identify the coefficients $\{\psi_H, \psi_H^l\}$ from the moments $[\Delta Z_{it}^p \cdot \Delta \epsilon_{it}^P] = 0$ where

$$\Delta Z_{it}^p = \left(B_{it,'00,3d}, B_{it,'00,3d} x_i^{land} \right).$$

The identifying assumption to obtain consistent estimates of ψ_H and ψ_H^l here is that unobservable shocks to house prices are not correlated with economic shocks based on national trends to the city, other than through the effect of the latter on housing demand and in-migration.

E.4.3 Industry choice elasticity a

The key estimating equation for the workers' industry choice elasticity parameter a is given by

$$\Delta \ln \pi_{is\iota t} = \alpha_{ist} + \alpha_{s\iota i} + a \Delta \ln W_{i\iota t} + u_{is\iota t},$$

where $\alpha_{ist} = \Delta \ln \left(\sum_{\iota=1}^{N_{ind}} W_{i\iota t}^a \epsilon_{s\iota i} \right)$. This means I can estimate a as the coefficient from a regression of changes in the local industry employment shares for each education group s on local log industry wage changes, controlling for a group-specific city trend α_{ist}. Note that $\Delta \ln \epsilon_{s\iota i} = 0$ has dropped out and that I have added a stochastic error term $u_{is\iota t}$ that captures any mismeasurement of wages.

From the model, national changes in industry demand represent shocks to labor demand that are plausibly exogenous with regard to local labor supply decisions. Therefore, I can use the national leave-one-out wage growth $\Delta \ln W_{\iota,-i,t}^{US}$. in industry ι as an instrument for local industry wage growth being higher relative to the city i average. The identification assumption for estimating a then becomes

$$E[\Delta \ln W_{\iota,-i,t}^{US} \cdot u_{is\iota t} | \alpha_{ist}, \alpha_{s\iota i}] = 0,$$

which would only be violated if industries that had higher national wage growth in *other* cities were systematically more likely to have higher or lower local employment shares through residual channels that are not operating through higher wages. Note that, by conditioning on city trends, I am controlling for any network migration effects or other city-level shocks. The elasticity a is therefore only identified off within-city differences in a given year between local industries. Moreover, in the empirical implementation I also condition on CZ-by-industry fixed effects $\alpha_{s\iota i}$ to control for long-run local industry trends that might be invalidating the exclusion restriction.

Note that the timing of the model implies that industry choices in period t are based on period t wages and are reflected in employment shares in period $t+1$, so I use employment shares in the next period to measure industry choice shares in period t.

F Data appendix

F.1 Migration network weight construction

This section discusses how the fixed migration exposure weights are estimated. In Section 3, I show that the migration exposure of city i to city j house prices is measured by

$$\psi^{ij} = \frac{1}{\bar{\psi}_i} \sum_{k \in N} \underbrace{\phi^{i \leftarrow k}}_{\substack{\text{In-migration} \\ \text{share for } i}} \underbrace{\mu^{k \rightarrow j}}_{\substack{\text{Out-migration} \\ \text{share for } k}},$$

where $\bar{\psi}_i = 1/\sum_{j:j\neq i}\sum_{k\in N}\phi^{i\leftarrow k}\mu^{k\rightarrow j}$ scales each city's relative migration exposures to sum to one. I use these migration exposure weights to construct the migration network house price changes, as well as the network instruments and any control variables measured at the migration network level.

One interpretation of this measure is that economic shocks to other cities can have effects on city i because of their *direct* migration connection, or due to *indirect* links as they share migration connections to another city in city i's migration network. This is analogous to the way that industry productivity shocks can indirectly affect other industries through shared local labor markets in network models in macroeconomics (Acemoglu et al., 2016).

There are several empirical issues to consider in constructing these migration weights: On the one hand, with a finite number of observations, actual year-to-year migration flows between cities are only a noisy measure of the true latent out- and in-migration probabilities. Therefore, μ^{ik}_{t-1} and $\phi^{i\leftarrow k}_{t-1}$ measured in a particular year might not capture true latent probabilities for city pairs with limited realizations in the data (Dingel and Tintelnot, 2020). In order to make these estimates of migration network connections more precise we would therefore want to average over several years.[107]

On the other hand, if migration is to some degree endogenous – as is one of the key arguments of this paper – then the migration network weights of other cities would fluctuate over time if they are measured on an ongoing basis. This would make any estimated effects difficult to interpret as they would combine changes in weights with changes in network prices. Moreover, if the effects of network price changes are auto-correlated, then migration network weights would anticipate future network price changes to some degree, leading to bias if price changes are not exogenous (Jaeger et al., 2018).

As an empirical compromise that addresses these concerns, I therefore hold the migration network weights fixed across years at baseline period values, measured using average migration flows over 1990-1995. The migration during this period, besides being the first years of migration data in the IRS migration flows sample, also has the advantage of preceding the dramatic house price boom-and-bust cycle that started in the mid-1990s, and is therefore unlikely to be caused by the house price dynamics that are of the greatest policy interest over the last three decades. Besides, holding the migration weights fixed at a baseline period level has a strong precedent in the literature using historical migration shift-share instruments (Altonji and Card, 1989; Boustan, 2010; Howard, 2020; Derenoncourt, 2019).

F.2 Bartik shift-share shock construction

Both the reduced form and the structural model estimation make use of a variety of different Bartik (1991) style shift-share instruments. This section summarizes the construction of all the different shift-share instruments used.

The fundamental idea behind these shift-share shocks is that, under some assumptions, one can obtain an instrument for exogenous local wage changes by using the interaction between local exposure to national industry trends with the size of those trends. Identification of the effect of local wage changes in such a setting follows either from the exogeneity of the cross-sectional variation in exposure (Goldsmith-Pinkham et al., 2018), or the exogeneity of national industry trends with regard to the exposure patterns (Borusyak et al., 2020).

In this paper, the model in Section 5 justifies the use of industry shift-share instruments because the industry demand side of the Armington (1969) style setup (see Equation 17) allows for national shocks in the form of either national shifts in consumption preferences towards industry ι (the γ_ι term in Equation 16), or arising from a common national component of industry productivity.

The model implies that I can construct log wage growth instruments by combining the exposure term in the form of the local wage bill share $\tilde{\omega}_{\iota,i,t_0}$ of workers in industry ι in a baseline period t_0 with shifters consisting of national wage growth $\Delta \ln W^{US}_{\iota,-i,t}$ in each industry ι in the form

$$B_{it,t_0} = \sum_{\iota} \tilde{\omega}_{\iota,i,t_0} \Delta \ln W^{US}_{\iota,-i,t}.$$

The industry averages of log wage growth are computed as leave-one-out measures to avoid mechanical correlation between the national trend estimate and city i wages (Borusyak et al., 2020). The source of industry data for all shift-share shock is the Quarterly Census of Employment and Wages (QCEW).

Baseline period. In order to minimize bias from endogeneity in the local industry exposure – which might result from auto-correlated national industry shocks (Goldsmith-Pinkham et al., 2018; Jaeger et al., 2018) – I fix

[107] Another approach to smoothing them statistically is shown in the construction of the migration-by-education-group data set in Appendix F.6.

industry exposure shares at the baseline level in some period t_0 for each analysis. This choice trades off efficiency of the instrument to prevent bias. In particular, the farther away the baseline period is from the years in the analysis, the less likely the baseline year exposure shares are to reflect the effect of national trends, because employment shares may have changed over time.

To ensure relevance of the instrument, I therefore adjust the baseline period for different analyses: For the reduced form estimation, the earliest period for which I have industry-level employment data – and also the first year of the panel used in my analysis – is 1990. The reduced form analysis thus uses 1990 industry employment shares for the full 1990-2017 panel of shift-share instruments. However, in the structural parameter estimation in Section sec:qea, I do not use any data from earlier than the year 2000. Therefore, I update the baseline period for the structural estimation and set the year 2000 as the baseline period, using year 2000 industry employment share shares for all shift-share instruments.

Industry aggregation level. The other data choice to be made in constructing the shift-share instruments is the level of industry aggregation to use in defining exposure and estimating wage trends. The trade-off involved in this decision is that using a more detailed industry definition for shift-share instruments (e.g. using 2-digit instead of 3-digit NAICS industries) will be less reflective of actual exogenous local wage shocks over time if local industry structure changes more between narrow categories than broad sectors, or if there are spillovers. For example, if the local presence of Credit Intermediation (NAICS 522) firms is linked to the local growth of Funds & Trusts (NAICS 525), then a shock to one of these narrow industries might also affect the other one. Moreover, firms or employees might shift between these industries. However, narrowly defined shift-share exposure would miss these effects. In addition, the QCEW data is more likely to suppress employment counts in narrow industries (for anonymity reasons) than in broader sectors, such that the estimate of actual industry structure might be noisier at the more detailed level. On the other hand, a broader industry definition might introduce noise due to the fact that growth in a particular 3-digit subsector at the national level might not affect a different subsector that is contained within the same 2-digit sector code. Aggregating across these subsectors might therefore mistakenly infer shocks to cities' local industries where there are none. As the exclusion restrictions for these two different aggregation levels are the same, I construct both measures. Where possible, I use the more detailed 3-digit aggregation, except for those application, such as the skill-specific shift-share instruments, that benefit from the more robust industry classification of the 2-digit measure.

Education group shift-share instruments. In addition to the shift-share instruments for overall changes in city wages, the model with heterogeneous skill groups (proxied by education levels) additionally suggests that we can construct skill-specific shift-shares of the form

$$B^s_{it,t_0} = \sum_{\iota} \pi_{\iota si,t_0} \Delta \ln W^{US}_{\iota,-i,t}.$$

The only difference in construction to the shift-share instrument for the city as a whole is that the weights on the national industry trends are now given by the exposure of group s to those trends, which is proxied by the share $\pi_{\iota si,t_0}$ of workers in group s who work in industry ι in the baseline year in city i.

Education shares by industry. Computing this employment share by education group requires the use of additional data on the education composition of different industries by city. I use Census microdata from IPUMS for the year 2000 on education and industry of employment of sampled individuals in each Adjusted CZ to compute the local education share by industry as follows: (1) Aggregate counts of employment by college / non-college education group and Census industry code (1990 definitions) and Adjusted CZs, from Census microdata sample for year 2000; (2) Create a probabilistic crosswalk of 1990 Census industry codes to 2-digit NAICS codes and apply it to the industry codes in the year 2000 microdata to crosswalk local employment by education into NAICS 2-digit industry codes; (3) Compute college education shares by NAICS 2-digit industry for each Adjusted CZ, replacing local value with national average for cities with small industry employment (¡200 workers); (4) Apply estimated education shares to QCEW local industry data (on employment and wage bills), using national average education share by industry, where no local Census data was available.

Then, I hold these education shares constant at their year 2000 level throughout. While this will reduce the power of the education shift-share instrument if education shares by industry shift substantially over time, constant education shares may reduce bias by a similar logic as used above to justify constant employment shares in the regular Bartik shift-share instruments. At the same time, microdata on education by industry is not continuously available after the year 2000 Census, which would make updating the skill shares impossible from a data perspective, if I wanted to use the estimation method described above.

Geographic aggregation. This paper uses two different geographic units, depending on the analysis. The

Figure A12: Wage trends by sector. The graph plots average annual wage growth by NAICS super sector. The data shown are average wages from the Quarterly Census of Employment and Wages, converted into real wages by adjusting for the urban consumer CPI, and indexed to 1990 values.

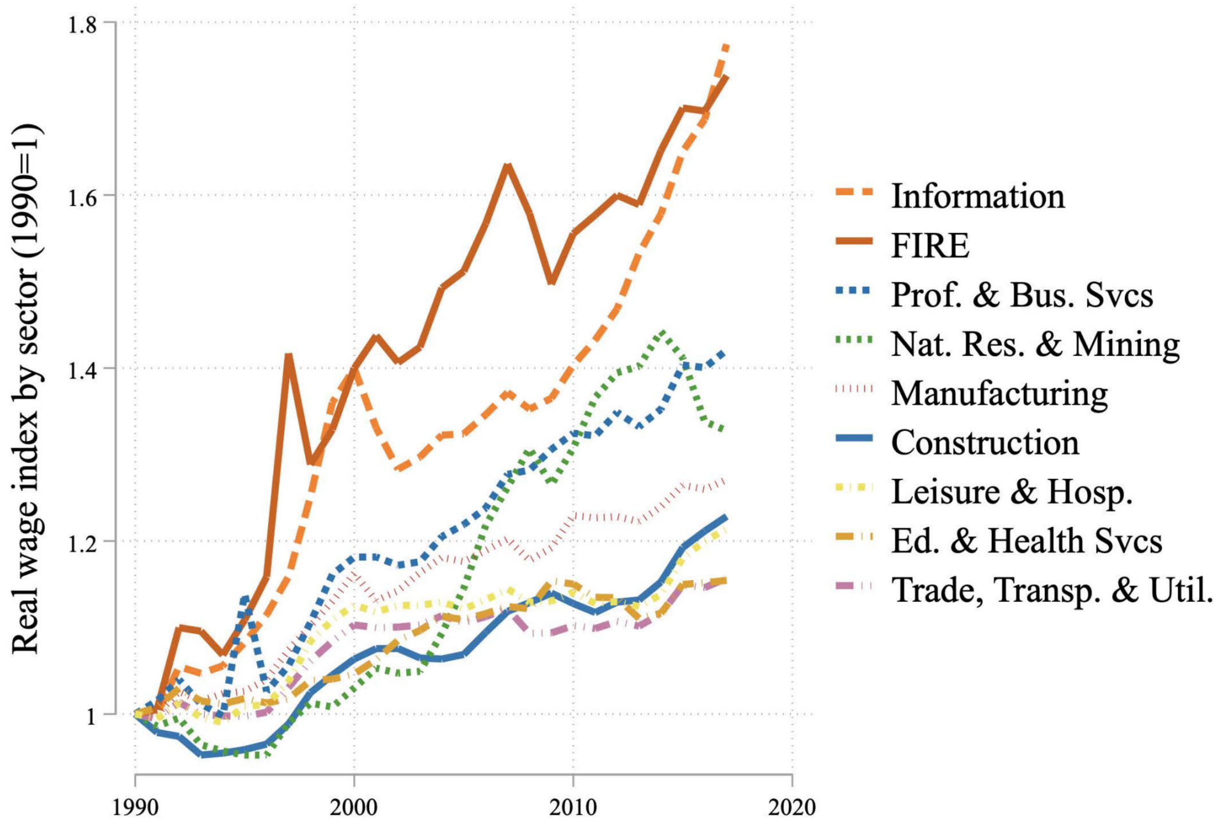

reduced form section focuses on the full set of continental U.S. commuting zones based on 1990 boundaries. In contrast, the structural estimation combines the commuting zones into "Adjusted CZs" that combine smaller CZs with their neighbours to ensure that all resulting unit have a population of at least 50,000. In each of these analyses, I compute shift-share instruments at the level of the geographic unit used. This means that the shift-share shocks in the reduced-form section are at the level of regular 1990 CZs, while the shift-share shocks in the structural model estimations are computed at the Adjusted CZ level.

Identifying variation. If there is heterogeneity in the causal effect of interest, the shift-share instruments, like any IV approach, will identify a local average treatment effect that represents a weighted average of the causal effects at different units. As is well known, the weights will depend on the degree to which the shift-share instruments represent larger shocks to some markets than others (Angrist et al., 2000). Therefore, it is of interest to understand the underlying identifying variation in the shifters used in the construction of these instruments. Three margins of variation are particularly relevant for my analysis: education, geography, and industry.

First, differences in wage trends across industries are driving the variation over time and space in the shift-share shocks in different CZs. These industry wage trends are shown for aggregated industry supersectors in the Appendix Figure A12. The graph shows that more high-skilled white collar sectors, such as information, finance, and business services have done comparatively well in recent decades, while manufacturing and construction, for instance, have fared less well in term of wages.

Second, the cross-sectional differences in city exposure to wage shocks, either directly or through their migration network come from the fact that cities vary in their industry structure. For instance, larger cities have benefitted more from the technological changes favoring skilled services industries (Eckert et al., 2019). Figure A13 illustrates this identifying variation by plotting the difference in wage trends by tercile of city size. As the left panel shows, wage growth has been higher in larger cities than in smaller cities, with larger cities experiencing double the wage

Figure A13: Wage trends by city size and education group. The graphs plot average annual wage growth by different city and education groups. The data shown are average wages from the Quarterly Census of Employment and Wages, converted into real wages by adjusting for the urban consumer CPI, and indexed to 1990 values. The subgroups in the left graph are terciles of Adjusted Commuting Zone population size, whereas the right panel shows college and non-college group wages for the top and bottom tercile by size. Wages by education group are computed by using year 2000 Census education shares by industry to weight industry wage growth.

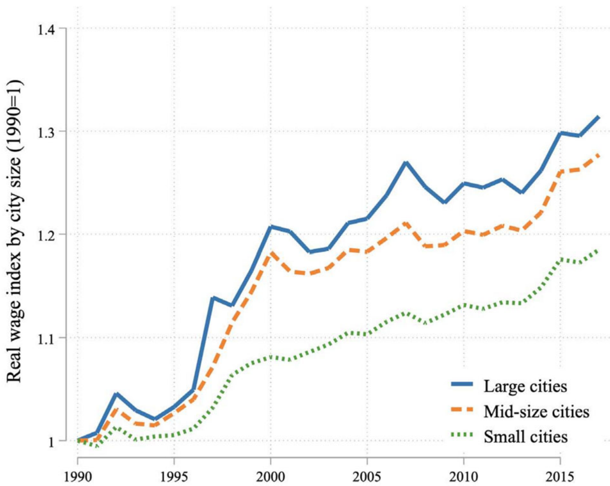

(a) By city size

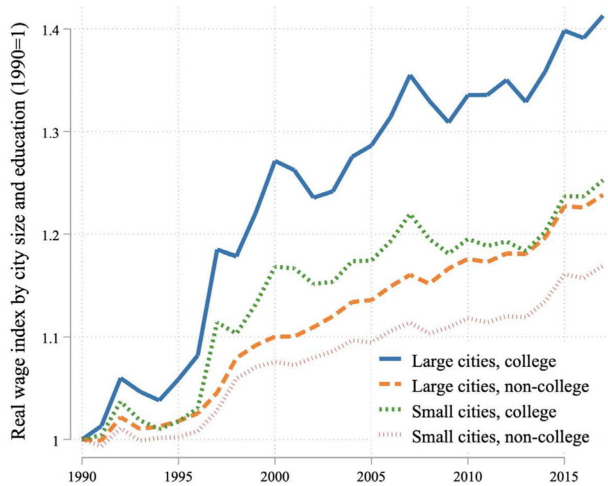

(b) By city size & worker education

growth of small cities during 1990-2005.

Third, displacement in cities requires variation across education groups in the wage growth that they experience. The right panel of Figure A13 shows that the differences in wage growth within city size terciles are also large. College workers in both large and small cities experience wage growth over the last three decades that was around 50% larger than the wage growth for non-college workers.

Overall, these graphs illustrate that the shift-share shocks based on differences in group and location exposure to national wage trends can generate substantial exogenous cross-sectional and time variation in wages, which I use to identify the parameters of interest in both the reduced form network analysis and the structural estimation.

Constructing house price instruments from wage shift shares. To construct shift-share shocks that can be used as instruments for local changes in house prices, I follow a large empirical literature (see, e.g. Diamond (2016)) in interacting wage shift share shocks with a proxy for exogenous housing supply constraints. As the measure of constraints, I use local land unavailability for construction x_i^{land} from Lutz and Sand (2019).[108] This measure captures geographic constraints to marginal housing construction, which would be expected to increase the slope of the housing supply curve, and thereby increase the responsiveness of house prices to the wage shocks. I aggregate county-level measures of land constraint into commuting zone average, weighting the county-level measure by the county population to reflect the likelihood that a city's residents are constrained by the geography.

Matching shift-share instruments to estimations. To summarize, note that wage shift-share instruments $B_{it,t_0}^{s,ind}$ can vary with the group s for which exposure shares are computed (full city, college, or non-college), the baseline year for the employment share exposure weights (year 1990 or 2000), and the industry aggregation (2-digit or 3-digit NAICS). While in most of my estimations all of these variants would be valid instruments under the same exclusion restrictions, weak instrument concerns necessitate using a subset where possible. I try to apply the following logic in choosing the appropriate instruments: (1) Long panel estimations that use data before 2000 use the 1990 weights, and shorter panels with only post-2000 data use the newer weights. (2) In industry-level regressions (the labor demand and industry choice parameter estimations), I match the aggregation level of the dependent variable and the shift-share instrument. (3) All education group shift-share instruments use 2-digit

[108] These are comparable to the Saiz (2010) land availability measures commonly used in the literature. Lutz and Sand (2019) build on his methodology to expand the number of covered cities and, among other things, improve the measurement of land availability for overlapping city areas and coastal locations.

industries because this increases the sample size and robustness of the estimated local education shares by industry.

F.3 Interest rate shock construction

As one example of spillover effects from a shock in Section 4.3, I construct predicted house price effects as a result of changes in financing costs (in the form of long-term interest rates) that interact with difference in local housing supply constraints.

I construct predicted house price spillovers from interest rate changes as follows: (1) I estimate the effect of national trends in long-term interest rates on house prices at the national level during the run-up to the housing boom of the 2000s, following a methodology similar to that in Glaeser et al. (2012). (2) I predict expected changes in house prices over the 1995-2007 period for each city, and aggregate these predictions into network terms, using the migration exposure weights. (3) I multiply this change in network house price by the estimated spillover coefficient from Column 6 in Table 4 and residualize this prediction with regard to the average values for each variable in the full set of controls used in my baseline analyses, as well as the direct effect of interest rate changes on the focal city. (4) I compare the residualized predicted house price spillover effect from interest rate changes to the actual house price growth over this period.

First, I estimate the effect of national trends in long-term interest rates on house prices at the national level during the run-up to the housing boom of the 2000s, following a methodology similar to that in Glaeser et al. (2012). That is, I run regressions of the form

$$\ln P_{it} = \alpha_i + \beta_1 Real10yrRate_t + \beta_2 Real10yrRate_t \times x_i^{\text{land}} + \epsilon_{it}$$

on data for 1990-2008. Here, I follow Glaeser et al. (2012) in computing the relevant long-term real interest rate as the differences between 10-year treasury rates and ten year inflation expectations obtained from the Livingston survey of inflation expectations.[109] Moreover, I include city fixed effects in the regressions to capture fundamental differences in house price growth, e.g. due to land constraints, that are not driven by interest rate changes.

Table A6: Interest rate effect on house prices

Dependent variable:	Log House Prices$_{i,t}$			
	(1)	(2)	(3)	(4)
Real 10-year rate	-0.13***	-0.13***	-0.13***	-0.05***
	(0.02)	(0.02)	(0.02)	(0.01)
Real 10-year rate× Land constraint		0.02		-0.25***
		(0.03)		(0.05)
Observations	13,338	12,996	13,338	12,996
CZ FE			✓	✓

Heteroskedasticity-robust standard errors clustered at the CZ level shown in parentheses: * p<0.10, ** p<0.05, *** p<0.01. Includes data from 721 CZs for 1990-2008.

The results of this estimation, weighting each data point by the city's population, are shown in Table A6. The interpretation of the magnitudes is that they represent the semi-elasticity of house prices with regard to a 1 ppt change in the real interest rate. For example, the estimate in Column 4 suggests that a 100 basis point decline in real rates is associated with a 5 log point increase in house prices in an unconstrained city, and a larger increase in more constrained cities.

I use the estimate with a land constraint interaction and fixed effects in Column 4 in order to construct predicted house price effects from interest rate changes for the run-up to the boom of the 2000s during 1995-2007. That is, I construct

$$\widehat{\Delta_{95\text{-}07} P_i^{ir}} = \hat{\beta}_1 \Delta_{95\text{-}07} Real10yrRate_t + \hat{\beta}_2 \Delta_{95\text{-}07} Real10yrRate_t \times x_i^{\text{land}},$$

where the change in the real rate from 1995 to 2007 in my data is −1.2 ppt. This decline in interest rates is shown in the times series graph in Figure A14, which also shows the concomitant rise in house prices over this period.

[109]This survey is done by the Philadelphia Fed and is available at: `https://www.philadelphiafed.org/research-and-data/real-time-center/livingston-survey`

Figure A14: Real long-term interest rates and house prices. The graph shows the time series of long-term real interest rates, calculated as the difference between 10-year treasury rates and ten year inflation expectations obtained from the Livingston survey of inflation expectations. For comparison, it also shows a population-weighted average across cities of repeat-sales house price indices from the FHFA, indexed to their level in 1991.

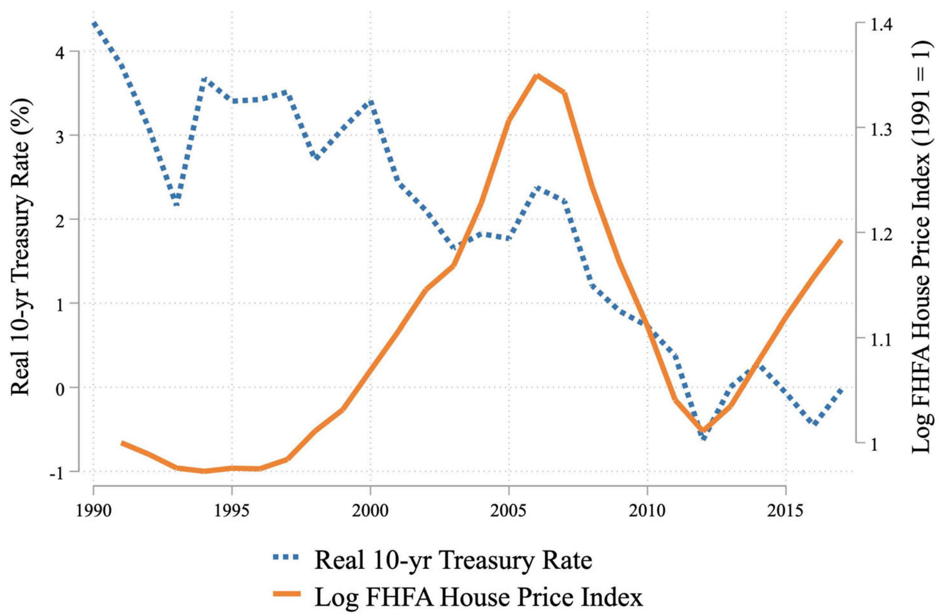

Note however, that the effect of interest here is the difference in cross-sectional exposure to this aggregate effect due to spillovers, not the size of the interest rate effect in the country as a whole. The graph also shows that the negative relationship between long-term real rates and house price breaks down during the Great Recession, so I focus on the period before 2008, where interest rates have been emphasized as one of the reasons for an increase in house prices.

I aggregate these predicted city-level effects into network terms, using the migration exposure weights. That is, I construct

$$\widehat{\Delta_{95\text{-}07}\mathcal{P}_i^{NW,ir}} = \sum_{j:j\neq i} \psi^{ij} \widehat{\Delta_{95\text{-}07}P_g^{ir}}$$

as the migration exposure to interest rate effects in other cities.

Next. I multiply this predicted change in network house prices for city i by the estimated spillover coefficient from Column 6 in Table 4 and residualize this prediction with regard to the average values for each variable in the full set of controls used in my baseline analyses, as well as the direct effect $\widehat{\Delta_{95\text{-}07}P_i^{ir}}$ of interest rate changes on the city i.

Last, I compare the residualized predicted house price spillover effect from interest rate changes to the actual house price growth over the 1995-2007 period, as well as each city's house price beta for that period, which results in the graphs shown in Figure 9.

F.4 Amenities index constructions

I follow Diamond (2016) and Almagro and Domínguez-Iino (2020) in measuring the availability of local amenities by using counts of local establishments in particular industries from the County Business Patterns data. These establishments are chosen based on their importance for serving tourists and residents for leisure activities, and therefore serve as proxies of how attractive amenities in a location are. Both Diamond (2016) and Almagro and Domínguez-Iino (2020) provide evidence that different demographic groups may value amenities differently. Supporting evidence that changes in local establishments are associated with changes in demographics is provided by Glaeser et al. (2018) who note that a change in the number of certain establishments such as wine bars can predict an increase in the average education level of a neighborhood.

Informed by these papers, I use establishment counts in the following NAICS codes as proxies for local amenities:

- Food Services and Drinking Places (NAICS 722)
- Grocery store industry (NAICS 4451)
- Motion picture theaters (NAICS 512131)
- Dry Cleaners (NAICS 8123)
- Clothing and Accessory Stores (NAICS 448)
- Museums, historical sites, and similar institutions (NAICS 7121)
- Sports teams (NAICS 71211)
- Scenic and Sightseeing Transportation (NAICS 487)

I use the crosswalks for different NAICS classifications from Eckert et al. (2020) together with the county-level annual CBP data for 1989-2016 to create a panel of establishments by county by NAICS 2012 code. Then, I aggregate the county-level data to 1990 commuting zones using the David Dorn crosswalk.

Moreover, for use in the regressions I scale the number of establishments in these categories by the IRS return-based population estimate for each CZ in order to express amenities in the form of a density of establishments per 1,000 residents.

The amenities index for each city is then computed as the first principal component $Amen_{it}^{1st}$ of the counts of establishments per 1,000 residents in the industries enumerated above that provide cultural or consumption services.

F.5 Combining smaller commuting zones into larger geographic units

As the pairwise migration data by education group that are used in the location choice regressions suffer from very small sample sizes if there are few data points per commuting zone, I improve the reliability of the estimates by combining 1990 commuting zones that have less than 50,000 residents with adjacent CZs until the combined area contains at least 50,000 residents.

I use the following algorithm to combine CZs: (1) Find all CZs with less than 50K residents; (2) For each of these CZs, find the smallest CZ within 50 miles centroid-to-centroid distance of any component county; (3) Combine each small CZ with the identified nearby target CZ. Repeat this loop until no more CZs have less than 50K residents. If no CZs within 50 miles exist to combine with, expand the search radius to 75 miles and then 100 miles. This adjustment mostly affects thinly populated areas in the middle of the U.S. Moreover, the adjustment also prevents many small CZs without data in particular years from dropping from the sample.

The resulting combination of CZs into larger "Adjusted CZs" is shown in the map in Appendix Figure A15, where adjacent CZ in the same color indicate CZs that are combined. Ultimately, the 741 original CZs are combined into 529 Adjusted CZs with the minimum population size, of which 512 Adjusted CZs are in the continental U.S. and form the main geographic units for the structural estimation.[110] For consistency, all the structural parameters are estimated using these Adjusted CZ geographic units, and these are also the units used in simulating the counterfactuals.

F.6 Migration flow by education group

In order to be able to estimate location choice parameters separately by education group, I need a data set of annual migration probabilities between city pairs for both college- and non-college-educated workers. Unfortunately, no single data set provides this information for U.S. cities over a long time horizon.

Given enough data, we could simply calculate the empirical choice probabilities for each state of interest where the state space consists of worker types, origins, destinations, and years. Unfortunately, the IRS data used in the reduced form estimation does not allow for distinguishing flows of different groups of workers. The best publicly

[110]This total excludes New Orleans, which is again excluded due to the large shock from Hurricane Katrina in 2005/2006.

Figure A15: Adjusted CZ definition based on combining CZs The map shows how small population CZs are combined with nearby CZs into Adjusted CZs with at least 50K residents. Adjacent CZs with the same color in the map (not grey) are combined into single units in the data.

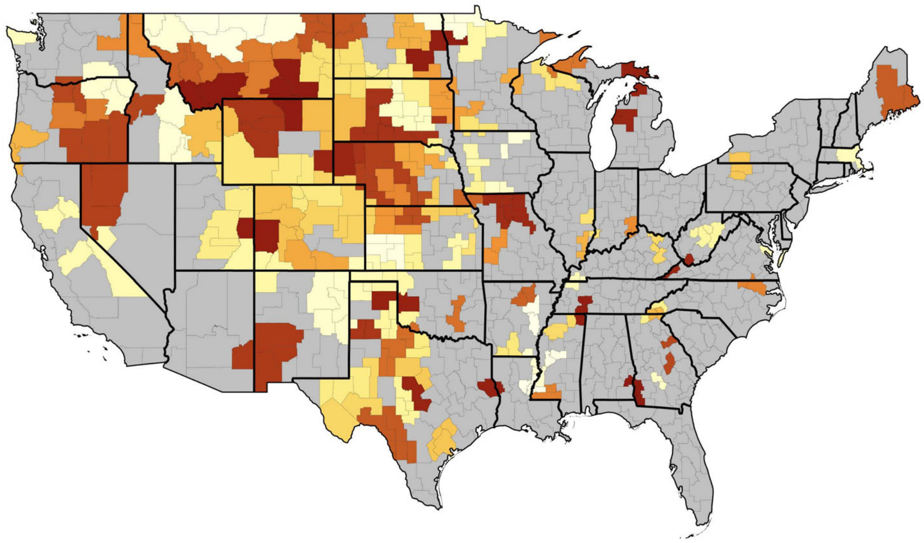

available migration data at an annual frequency for the U.S. that also contains migrant characteristics, is from the American Community Survey (ACS) for 2005-2017, with sample sizes in the single-digit millions for most years. However, even in the most basic setup, the state variables consist of all permutations of high and low worker skill with potential origin continental U.S. commuting zones and destination commuting zones, so up to $2 \times 512 \times 512 = 524,288$ different states in each year. As a result, the raw estimates of CZ-to-CZ migration flows by education are noisy because some cells are necessarily estimated with a small sample size.[111] To improve the precision of my estimates, I therefore apply statistical techniques for data smoothing and imputation that allow me to make use of additional information contained in the ACS data, combine information across units and years, and incorporate information from the IRS migration data.

Therefore, in order to overcome the limitations of individual data sets, I decompose the migration probabilities by education group into components that can be reliably estimated and employ statistical tools to mitigate weaknesses in individual data sets, for instance by combining information from multiple sources.

In particular, I decompose the share of people in education group s in city i moving to city k as follows:

$$\mu_{st}^{ik} = \underbrace{\frac{m_{st}^{ik}}{\sum_s m_{st}^{ik}}}_{\text{Skill } s \text{ share in } i \to k \text{ flows}} \cdot \underbrace{\mu_{out,t}^{ik}}_{\text{Share of } i \to k \text{ in all out-mig.}} \cdot \underbrace{(1-\mu_t^{ii})}_{\text{Out-mig. rate from } i} \cdot \underbrace{\frac{L_{i,t-1}}{L_{is,t-1}}}_{1 \text{ / Pop. share of type } s} \tag{42}$$

where $m_{st}^{ik} = \mu_{st}^{ik} L_{is,t-1}$ are the number of migrants of type s moving from i to k in period t. As noted above, I estimate these components separately, combining data from both the ACS and the IRS.

The construction of my data set of flows by education between commuting zones proceeds in several steps: (1) Imputing CZ-to-CZ total migration flows $\mu_{out,t}^{ik}$ from IRS data. (2) Predicting education shares for flows. (3) Smoothing gross migration rates $(1-\mu_t^{ii})$ to account for IRS data issues. (4) Combining the components of migration rates into predicted pairwise migration probabilities.

Imputing CZ-to-CZ flows from IRS migration data. In general, I use the aggregate data from the

[111]While the effective number of states is substantially smaller as the migration flow matrices are sparse, a more elaborate state space, for instance taking into account migrant age or state of birth, would expand the number of states to estimate even further. Moreover as migrants only represent a small share of the surveyed individuals in each year, the effective sample size for migration purposes is much smaller.

Internal Revenue Service Statistics of Income (IRS SOI)[112] based on the near-universe of tax returns as a measure of the relative size of overall migration flows between locations. The IRS counts moves in the form of changing adresses on tax returns and also totals up exemptions claimed on moving tax returns – which I will use as a proxy for the number of people moving.

The IRS data comes in the form of county-to-county flows, as well as total flows in and out of each county and flows to and from some aggregate regions. A peculiar feature of the IRS data is that it only records county-to-county flows by name when the flows add up to at least 10 tax returns for the pair. Smaller "unnamed" flows are, however, included in the out- and inflow totals as well as in more aggregate flow totals by census region.

In order to avoid county pairs dropping in and out of the data if they become too small and to capture the fact that true migration probabilities are unlikely to be precisely zero (Dingel and Tintelnot, 2020), I impute small flow probabilities by allocating unnamed outflows from each county using a hierarchical "empirical Bayes" approach. I assume that the likelihood of flows between regions and to cities within regions can be described as a multinomial distribution, the parameters of which come from a Dirichlet distribution. Then, I use observed aggregate information on regional flows, and observed named flows as empirical estimates of the prior for flow shares going to each city. These priors are then used to allocate unnamed flows to cities. Last, allocated unnamed flows and named flows are smoothed using again a Multinomial-Dirichlet model to estimate non-zero latent probabilities of pairwise migration between cities.

Allocating unnamed outflows to destination regions. To allocate the unnamed outflows N_c^{out} from each county c, I follow the following process: First, I subtract any foreign flows from the unnamed flows to obtain domestic unnamed flows. Second, I initially assign these unnamed flows to the regional totals listed as going to the same state or out-of-state to one of the four Census regions (Northeast, Midwest, South, West). Any remainder is assigned equally to each of these 5 categories. Third, I smooth the allocation of unnamed flows by assuming that the distribution across the 5 destination categories follows a multinomial distribution, with probabilities θ_{ic}^{reg} of an unnamed flow from c going to a destination in each of these categories i. That is, I assume that the probabilities θ_{ic}^{reg} of an unnamed flow going to a destination in each of these categories have a joint density

$$p(\theta_{1c}^{reg}, \ldots, \theta_{5c}^{reg} | \boldsymbol{\alpha}) \propto \prod_{i \in \{1,5\}} (\theta_{ic}^{reg})^{\alpha_i - 1},$$

where $\sum_{i \in N} \theta_{ic} = 1$ and the parameters α_i characterize the prior distribution. I choose a non-informative prior of $\alpha_i = 1 \; \forall i$ (a uniform distribution over all possible values of θ_{ic}^{reg}), and treat the allocated unnamed migration flows to each region as data $\boldsymbol{y}$ that is used to update the parameter estimate (Gelman et al., 2013). The marginal means of the posterior distribution will be given by

$$E[\theta_{ic}^{reg} | \boldsymbol{y}] = \frac{n_{c \to i} + \alpha_i}{N_c^{out} + \sum_i \alpha_i},$$

where the $n_{c \to i}$ are the unnamed flows initially allocated to region i. The posterior estimate of the unnamed flows to each region is then given by $\hat{n}_{c \to i}^{reg} = E[\theta_{ic}^{reg} | \boldsymbol{y}] \cdot N_c^{out}$. I compute $\hat{n}_{c \to i}^{reg}$ for each county in each year.

Allocating unnamed outflows to counties within regions. Once unnamed outflows are allocated to regions, I distribute them across counties in each region. For each region, unnamed flows can only be allocated to counties that do *not* already have named outflows registered in the IRS data. Again, I assume that flows to each county within a region follow a multinomial distribution. For each of those potential destination counties within a region, I again form a non-informative prior with a uniform density over migration probabilities, i.e. $\alpha_k = 1 \; \forall k$ for the share of flows to that region that would go that county. The data used to update that prior is the share of all *named* flows to any county in that region – from other counties in the set C_{all} – that go to each county. That is, I use the data of observed named flows to update my prior for allocating unobserved flows, which can be thought of as an approximation to a complete hierarchical Bayesian analysis (Gelman et al., 2013).

To be precise, the Dirichlet posterior mean of the share θ_{kic}^{cty} of outflows from c going to county k, conditional on going to region i is

$$E[\theta_{kic}^{cty} | \boldsymbol{y}] = \frac{\sum_{j \in C} n_{j \to k} + \alpha_k}{\sum_{k \in C_{ic}^{un}} \sum_{j \in C_{all}} n_{j \to k} + \sum_{j \in C_{ic}^{un}} \alpha_j},$$

where C_i^{un} is the set of counties that are potential unnamed destinations in region i from the perspective of origin county c. Then, final imputed unnamed flows from county c to county k in region i are given by $\hat{n}_{c \to k} = \hat{n}_{c \to i}^{reg} \cdot E[\theta_{kic}^{cty} | \boldsymbol{y}]$, calculated separately for each year.

[112] Available at URL: `https://www.irs.gov/statistics/soi-tax-stats-migration-data`

Last, I apply David Dorn's crosswalks to aggregate county-to-county flows to 1990 commuting zones (CZ), combining small CZs into entities of at least 50K residents (see Section F.5), and treating within-CZ flows as non-migrants. Then, I drop any flows from and to CZs in Hawaii and Alaska, focusing on cities in the contiguous U.S., as well as flows to and from New Orleans, which experiences highly abnormal migration flows as a result of Hurricane Katrina in 2005/2006. For the remaining CZs, I aggregate all observed or imputed outflows to the other CZs in the sample and compute the share of outflows going to each destination CZs, which gives me the conditional migration probability $\mu^{ik}_{out,t}$.

Gross migration rates. The previous imputation was for migration shares conditional on leaving the CZ. I also use the IRS data to compute the overall level of out-migration $1 - \mu^{ii}_t$ from each CZ in each year. However, the IRS changed its methodology for inferring moves from tax returns starting with the tax returns for income year 2011, generating a break in the series of gross migration rates (Molloy and Smith, 2019). After this transition, gross migration rates in the IRS data seem to exhibit a larger-than-normal year-to-year volatility that is not supported by other data sources (e.g. the CPS migration rates). However, there is no reason to believe that this change in gross migration rates affects *relative* migration shares for different destinations and differences between CZs in migration activity systematically.

To remove this extraneous volatility in gross migration in later years from the data, I impute the overall level of gross migration for each CZ in the following way: I regress aggregate population-weighted averages of gross inter-city outmigration from the IRS ($Mig_{IRS,t}$) on the equivalent measure of gross migration to the continental U.S. in the American Community Survey ($Mig_{ACS,t}$) in a regression of the form

$$Mig_{IRS,t} = \alpha + \beta Mig_{ACS,t} + \epsilon_t$$

for the period 2005-2010. I predict gross IRS migration for 2011-2017 from the observed ACS values and the predicted relationship from the regression. Then, I impute individual CZ gross migration rates by applying the observed ratio of their gross out-migration to the average in the IRS data to the new imputed average value for post-2011 data. That is, I compute

$$\widehat{1 - \mu^{ii}_t} = \widehat{Mig}_{i,t} = \frac{Mig_{i,IRS,t}}{Mig_{IRS,t}} \left(\hat{\alpha} + \hat{\beta} Mig_{ACS,t}\right),$$

which is my measure of overall CZ outmigration rates. While I could, in theory, rely only on the ACS data to construct the same flows, the ACS data is meant to capture relative representation of population at the local level, but is not designed to provide an accurate count of population (and therefore migration) *levels*. The IRS data, which uses the universe of tax returns therefore seems better suited to estimate the size of total flows between cities. The ACS data is used instead to estimate the *composition* of flows and local populations.

Empirical Bayes shrinkage estimate of education flow shares. Next, I compute the share of the different education groups in flows between CZs. Because the ACS data on city-to-city flows is very noisy on an annual basis, so is the share of these flows that can be attributed to different education groups. However, due to the fact that migration by skill share needs to fulfill adding-up constraints, I can use information from other parts of the ACS data to predict pairwise flows and then apply an empirical Bayes shrinkage estimator to combine these predicted skill shares with the actually observed values.

I proceed in three steps: First, I use a version of the post-LASSO estimator of Belloni et al. (2013) to predict education shares for city-to-city flows. This involves first applying a Least Absolute Shrinkage and Selection Operator (LASSO) to select the origin and destination city characteristics that best predict the observed non-college share nc_i of migrants between two cities. This estimator chooses the coefficients on predictive characteristics by solving:

$$\hat{\beta}^{\text{Lasso}} = \arg\min_{\beta} \sum_{i}^{n} \left(nc_i - \sum_{j}^{p} x_{i,j} b_j \right)^2 + \lambda \sum_{j}^{p} \|b_j\|$$

where i indexes a CZ pair-year unit, and $x_{i,j}$ represents the candidate characteristics for non-college city pair flow share prediction, which consist of: the non-college share among all inflows into the origin city and destination city; all migrant outflows from the origin city; the non-college share of the total population in the origin and destination cities; the log of the total and non-college population levels in the origin and destination cities; and a constant. The penalty term λ is chosen by cross-validation with 10 folds, and the estimator applied to the pooled city pair data for 2005-2017 selects a non-zero coefficient for 6 out of the 10 candidate variables, omitting the total population levels, the log of the destination non-college populations, and the non-college population share in the origin city.

The selected non-zero coefficients with non-zero coefficients are then used in an OLS forecasting regression, and the OLS coefficients $\hat{\beta}^{\text{OLS}}$ are used to forecast a predicted non-college share

$$\hat{nc}_i^{pred} = X_i \hat{\beta}^{\text{OLS}},$$

and compute the standard error of the non-college share forecast s_i^{pred}. To reduce outliers generated by the linear forecast, I cap predictions to be no bigger than 99%.

Second, I compute the estimated non-college shares of city-to-city flows $\hat{nc}_i^{ACS}$ in the ACS microdata. The standard error of the ACS non-college share $\hat{nc}_i^{ACS}$ in flows computed from IPUMS microdata (Ruggles et al., 2020) can be computed approximately as (Bureau, 2005–2017):

$$s_i^{ACS} = df \cdot \sqrt{\frac{99}{\hat{n}_i^{ACS}} \hat{nc}_i^{ACS}(1 - \hat{nc}_i^{ACS})},$$

where $\hat{n}_i^{ACS}$ are the total flows observed for city pair-year i in the ACS data – the denominator of $\hat{nc}_i^{ACS}$ – and df is a design factor that reflects the ACS sample design and is obtained for each year from Bureau (2005–2017).[113] Note that for non-college shares close to zero or one, I follow the guidance in Bureau (2005–2017) and substitute 2% and 98%, respectively, for the purpose of the standard error calculation.

Third, I combine the predicted value $\hat{nc}_i^{pred}$ with the noisy observed non-college share $\hat{nc}_i^{ACS}$ for each city pair in each year by taking a weighted average, with weights that account for the relative uncertainty of the raw ACS estimate and the predicted value. To combine the two estimates, we can treat the predicted value as a prior with normal distribution $\mathbb{N}(\hat{nc}_i^{pred}, s_i^{pred})$, and the ACS data as being generated by a process with a normal distribution $\mathbb{N}(\hat{nc}_i^{ACS}, s_i^{ACS})$. Then, for any symmetric loss function, the optimal Bayes posterior estimator for the non-college share is

$$\hat{nc}_i^{eb} = \left(\frac{s_i^{ACS}}{s_i^{ACS} + s_i^{pred}}\right) \hat{nc}_i^{pred} + \left(\frac{s_i^{pred}}{s_i^{ACS} + s_i^{pred}}\right) \hat{nc}_i^{ACS}.$$

Intuitively, this "empirical Bayes" estimator $\hat{nc}_i^{eb}$ adjusts the raw non-college shares by moving them towards their expected value – "shrinking" the deviation — and does so to a greater degree if the raw estimate was based on a smaller sample size, and thus has a greater standard error. This approach to noise reduction by combining two estimates with one being treated as a Bayesian prior, although it is empirically constructed, is often called "empirical Bayes shrinkage" and has had a long history of statistical applications (Morris, 1983; DuMouchel and Harris, 1983; Gelman et al., 2013). One advantage of this method is that I can impute skill shares even for migration city pairs that are not observed in the ACS data, but for which the IRS data records migration flows. For values missing in the ACS data, I assume $s_i^{ACS} \to \infty$, such that the estimator loads entirely on the predicted value $\hat{nc}_i^{pred}$ in the estimation. As a result, this method yields an estimate of the non-college share - and therefore also the college share - of the migrant flows between each city pair and for each year. I use these estimates of $\frac{m_{st}^{ik}}{\sum_s m_{st}^{ik}}$ in equation 42 to compute the city pair migration probabilities by education group.

As this approach to data cleaning is "statistical" rather than fully model-dependent, there might be a concern that the smoothing and imputation procedure removes variation of interest from the data. However, for large pairwise migration connections, the effect is minimal as the empirical observation dominates the weak prior in generating migration probability estimates. For the city pairs with zero observations where the imputation procedure and allocation of unnamed flows makes relatively large changes relative to the observed flows, we know that the "zero" flow estimate is wrong for some city pairs from the discrepancy between total and named flows. Assuming these flows to be precisely zero would therefore be equivalent to discarding this information. This issue is exacerbated when taking logs of the flows: a common approach in the applied literature to this "zero flow" problem is to discard the zero observations implicitly when taking logs, use ad hoc adjustments such as adding 1 to each observation, or use nonlinear estimators that treat the observations as precisely zero. By using information available in other parts of the data set to determine which of the censored observations are more or less likely to actually be zero, I therefore think that I am adding information relative to these alternatives.

Substituting all the estimates of migration flow components described above into Equation 42 then yields the estimate of education-share specific location choice probabilities $\widehat{\mu}_{st}^{ik}$ that I use in the structural estimation.

[113] The design factors obtained from Bureau (2005–2017) for the applicable characteristic "Residence 1 Year Ago" is 3.0 for 2005–2008, 2.9 for 2009–2011, and 2.8 for 2012–2017.

F.7 Quality-adjusted house price index

In order to be able to compare house prices across cities, it is necessary to adjust them for differences in the quality of the housing stock. I follow Albouy and Ehrlich (2018) in constructing a quality-adjusted index of commuting zone-level house prices for the year 2000. House price data are obtained from the United States Census Integrated Public-Use Microdata Series (IPUMS), from Ruggles et al. (2020).

House price indices for each commuting zone j are calculated from a 5% sample from year 2000 Census. The sample is restricted to owner-occupied units. I regress the logarithm of house value $\ln P_{ij}$ for each household i on hedonic control variables X_{ij}, and indicator variables for each CZ cell. The regression specification is

$$\ln P_{ij} = \beta' X_{ij} + \psi_j + \epsilon_i,$$

where the estimated CZ fixed effects ψ_j are then treated as the commuting zone-level house price premia or discounts.

The housing characteristics X_{ij} included in the regression are:

- 10 indicators of number of units in the building
- 9 indicators for when the building was built (by decades or 5-year spans)
- 9 indicators for the number of rooms, 6 indicators for the number of bedrooms, as well as indicators for number of rooms interacted with number of bedrooms
- 2 indicators for complete plumbing and kitchen facilities

I follow Albouy & Ehrlich (2018) in estimating the regression in two stages: First, the regression is run weighting by census-housing weights, adjusted by the weight of the PUMA in the commuting zone (following David Dorn's crosswalks). A new value-adjusted weight is calculated by multiplying the CZ-adjusted census-housing weights by the predicted house value from this first regression using housing characteristics alone, but omitting CZ differences. A second regression is run using these new weights on the housing characteristics, along with the CZ indicators. The housing-price indices are obtained from the CZ fixed effect variables estimated in this second regression.

In order to compare house prices across cities, I then compute the predicted price of a "standard" home in each city, where "standard" is defined as using the mode of each variable in the hedonic housing characteristics. The baseline price is thus calculated for a single-family detached home, built in the 1990s, with 3 bedrooms and 6 rooms overall, and with complete kitchen and plumbing facilities. To this baseline home value is added the commuting zone house price indicator to compute quality-adjusted prices in each commuting zone.

This index is the converted into a panel under the assumption that quality differences are constant at their year 2000 values by constructing CZ-level values for other years using the FHFA repeat-sales index. That is, I apply the local house price growth in the FHFA to the year 2000 house price index to construct a 1990-2017 panel of house prices that has been adjusted for year 2000 quality differences.

F.8 Real wage index calculation

In order to compute the levels of real wages across cities for the welfare comparison, I need to make some assumptions about the distribution of industry productivity $\epsilon_{s\iota i}$. In particular, I assume that the attractiveness of retail trade (NAICS codes 44-45) does not vary across cities and is normalized to one.[114] That is, $\epsilon_{si,\text{Retail}} = \epsilon_{sj,\text{Retail}} = 1 \; \forall i, j \in N$. This allows me to use the employment shares and industry-level wages to infer the relative productivity terms for all other local industries from

$$\frac{\pi_{ist,\iota}}{\pi_{ist,\text{Retail}}} = \frac{W^a_{i\iota t}\epsilon_{s\iota i}}{W^a_{it,\text{Retail}}}$$

which I can then aggregate into a wage index

$$\widetilde{W}_{ist} = \Gamma\left(\frac{a-1}{a}\right)\left(\sum_{\iota=1}^{N_{ind}} W^a_{i\iota t}\epsilon_{s\iota i}\right)^{\frac{1}{a}}.$$

[114]This can be justified by noting that the employment share of retail is relatively large in all cities and varies little relative to that of other industries.

In order to convert this into the real wage index used in the welfare calculation, I adjust the wage index for differences in house prices:

$$R^w_{ist} = \frac{\widetilde{W}_{ist}}{P_{it}^{\alpha_s}}$$

The quality-adjusted house price levels are computed as described in Appendix Section F.7. Without loss of generality for the welfare calculations, all real wages are expressed relative to those for the New York City commuting zone.

This real wage index is computed for observed reference years (years 2000 and 2012), and are then converted into baseline steady state level differences by applying the relative change in real wages along the steady state path to the relative real wage index.

Note also that the log change in the income index can be written as

$$\Delta \ln \widetilde{W}_{ist} = \frac{1}{a} \ln \left(\sum_{\iota=1}^{\mathcal{I}} \pi_{is\iota t} \frac{W^a_{i\iota t}}{W^a_{i\iota,t-1}} \right).$$

That is, the growth in expected income for a worker of group s consists of a weighted average of the growth in industry wage rates, with the weights depending on the expected suitability of those industries for the worker – which is reflected in past workers' industry choices.

F.9 Trade flow link construction

One alternative measure of city-to-city links that I consider is the value of trade flows between them. This section describes how I construct a measure of the trade flow link between cities.

Trade flow data source. As the basis for the computation of trade flows between areas, I use the Commodity Flow Survey (CFS) Public Use Microdata for 2012 provided by the U.S. Census Bureau. This data is provided at the level of CFS areas, which are aggregations of counties, and captures the value and weight of individual shipments between CFS areas.

Computing CFS area trade links. First, I aggregate the microdata into origin-destination CFS area pair shipment values by summing across individual items, adjusting for a given weight factor that captures the representativeness of each line item. Then, for each origin CFS area, I compute the share of total trade value sent to each other CFS area. This is the measure of the strength of trade links between areas, which I then crosswalk to commuting zones.

Area crosswalk to commuting zones. I use a crosswalk from CFS areas to counties and from counties to 1990 commuting zones to probabilistically map CFS areas to commuting zones. CFS areas tend to be larger than commuting zones: the continental U.S. data contains 70 CFS areas, which I map into 169 commuting zones, and only 8 CZs map into more than one CFS area. I assume that CZs that are fully contained within a CFS area inherit all of the trade links of that CFS area. That is, their links with other CFS areas and the CZs contained within them are the same as for the CFS area that they are a part of. Where a CZ contains counties that are part of different CFS areas, I take a weighted average of the links of those CFS areas with other areas, weighting them by the share of the respective counties in the CZ's total population in the year 2000. The result of this crosswalk is a CZ-to-CZ measure of the *relative* strength of trade links with other CZs from the perspective of a CZ as the origin of trade flows.

G Counterfactual estimation algorithms

G.1 Rewriting the dynamic spatial equilibrium model in changes

The goal of this section is to rewrite the dynamic model in changes, such that it can then be simulated based on series of growth rates without having to know the levels of unobservable fundamentals. This derivation builds on the method used in Caliendo et al. (2019), which I adapt to the case of heterogeneous migrant groups.

I start by taking the ratio of migration probabilities from equation 14 between two time periods to obtain

$$\mu_{s,t+1}^{ik} = \mu_{st}^{ik} \cdot \frac{\exp(\upsilon_{t+1}^{ks} - \upsilon_t^{ks})^{\frac{1}{\theta}}}{\sum_j^N \frac{\exp(\upsilon_{t+1}^{js} - \tau_s^{ij})^{\frac{1}{\theta}}}{\sum_j^N \exp(\upsilon_t^{js} - \tau_s^{ij})^{\frac{1}{\theta}}}}$$

$$= \frac{\mu_{st}^{ik}(\dot{\mathcal{U}}_{t+1}^{ks})^{\frac{1}{\theta}}}{\sum_j^N \mu_{st}^{ij}(\dot{\mathcal{U}}_{t+1}^{js})^{\frac{1}{\theta}}} \tag{43}$$

Here, I have made use of the notation $\mathcal{U}_t^{is} = \exp(\upsilon_t^{is})$, and $\dot{x}_{t+1} = \frac{x_{t+1}}{x_t}$. Also, note that I have been able to eliminate moving costs from the expression because they are assumed to be constant over time.

Next, I difference the conditional value function as stated in equation 12 to obtain

$$\upsilon_{t+1}^{is} - \upsilon_t^{is} = U_{is,t+1} - U_{ist} + \beta E[V_{t+2}^{is} - V_{t+1}^{is}]$$

Taking exponentials of both side, using equation 13 to first substitute for the option value of location i, and then substituting the definition of flow utility U_{ist} and the definition of $\mathcal{U}_t^{is}$, we get

$$\dot{\mathcal{U}}_{t+1}^{is} = \left(\frac{A_{is,t+1}}{A_{ist}}\right)\left(\frac{\widetilde{W}_{is,t+1}}{\widetilde{W}_{ist}}\right)\left(\frac{Q_{i,t+1}}{Q_{it}}\right)^{-\alpha_s}\left(\frac{\sum_k^N \exp(\upsilon_{t+2}^{ks} - \tau_s^{ik})^{\frac{1}{\theta}}}{\sum_k^N \exp(\upsilon_{t+1}^{ks} - \tau_s^{ik})^{\frac{1}{\theta}}}\right)^{\beta\theta}$$

$$= \dot{A}_{is,t+1}\dot{\widetilde{W}}_{is,t+1}(\dot{Q}_{i,t+1})^{-\alpha_s}\left(\sum_k^N \mu_{s,t+1}^{ik}(\dot{\mathcal{U}}_{t+2}^{ks})^{\frac{1}{\theta}}\right)^{\beta\theta} \tag{44}$$

Here, I have also assumed rational expectations on the part of the migrants to drop the expectation operators for future value terms.

Assuming no extrapolation, the wage index can be written in changes based on the definition in equation 10 as

$$\dot{\widetilde{W}}_{is,t+1} = \left(\sum_\iota \pi_{\iota ist}\left(\dot{W}_{\iota i,t+1}\right)^a\right)^{\frac{1}{a}}$$

$$= \left(\sum_\iota \pi_{\iota ist}\left((\dot{L}_{i,\iota,t+1})^{\eta_{LD}}(\dot{\tilde{D}}_{i\iota,t+1})^{\frac{1}{\sigma}}\right)^a\right)^{\frac{1}{a}}, \tag{45}$$

where in the second line I have substituted from the labor demand equation 17 written in changes as $\dot{W}_{\iota i,t+1} = (\dot{L}_{i,\iota,t+1})^{\eta_{LD}}(\dot{\tilde{D}}_{i\iota,t+1})^{\frac{1}{\sigma}}$.

Similarly housing cost changes can be rewritten starting from their definitions in equations 10 and 20:

$$\dot{Q}_{is,t+1} = \dot{P}_{is,t+1}e^{-\Delta\alpha_t}$$

$$= \dot{\tilde{\phi}}_{i,t+1}e^{-\Delta\alpha_t}(\dot{HD}_{it})^{\tilde{\phi}_i^H}, \tag{46}$$

where $\dot{HD}_{it} = \frac{\sum_\iota \sum_s \alpha_s W_{\iota i,t+1} L_{is,\iota,t+1}}{\sum_\iota \sum_s \alpha_s W_{\iota i,t} L_{is,\iota,t}}$.

The evolution of local populations by education group over time is simply restated here as

$$L_{i,t+1} = \sum_s L_{is,t+1} = \sum_s \sum_k \mu_{st}^{ki} L_{kst}. \tag{47}$$

Last, to obtain the change in local industry employment, we begin by computing changes in industry choice shares by education group, which can be done analogously to migration shares to obtain

$$\dot{\pi}_{is\iota t} = \frac{\pi_{is\iota,t-1}\dot{W}_{i\iota t}^a}{\sum_{\iota=1}^{N_{ind}} \pi_{is\iota,t-1}\dot{W}_{i\iota t}^a \epsilon_{s\iota i}}$$

Then, industry employment changes are given by

$$\dot{L}_{is\iota,t+1} = \dot{\pi}_{is\iota t}\dot{L}_{is,t+1} \tag{48}$$

Note that equations 43, 44, 45, 46, 47, and 48 define a system of equations that can be solved for changes in population, house prices and wages, starting from a particular period's values as long as we assume that $\dot{\tilde{\phi}}_{i,t+1}\dot{\tilde{\phi}}_{t+1}e^{-\Delta\alpha_t} = 1 \quad \forall t$ for the simulation, and we are given a convergent series of changes in fundamentals $\{\{\{\dot{A}_{is,t+1}, \dot{\tilde{D}}_{i\iota,t+1}\}_{s=1}^{\mathcal{S}}\}_{i=1}^{N}\}_{t=0}^{\infty}$ where $\lim_{t\to\infty}(\dot{A}_{ist}, \dot{\tilde{D}}_{i\iota,t+1}) = (1,1)$ and observe baseline period values $(L_{ist}, W_{i\iota st}, Q_{it}) \ \forall(s,i,\iota)$.

G.2 Computing stationary steady-state equilibria

This section details the algorithm uses to compute stationary steady states for a given year's observables under the assumption of no further changes in fundamentals.

The algorithm starts by choosing some sufficiently large time period $t+T$, e.g. $T = 100$, at which the steady state should have been reached. Then, we need to initialize a vector of utility growth $\dot{\mathcal{U}}_{t+2}^{ks}$ for each city that is set to 1 at time T and some reasonable starting value (e.g. one) at all other points in time. Then, run the following loop until the path of populations and utility changes converges between iterations:

1. Initialize city population by group s in period $t+1$ as $L_{is,t+1} = \ell_{i,s,t+1} + \sum_j \mu_{st}^{ji} L_{jst}$. Here, ℓ_{it} are exogenous aggregate population changes, i.e. population changes not accounted for by domestic migration – which I compute from the change in total IRS exemptions, assuming that the exogenous population change has the same education distribution as the city itself in period t. For the steady-state computation, I assume that the average exogenous population change of the previous 5 years continues for another 5 years and then goes to zero. Then, use this together with the baseline population levels to compute $\dot{L}_{is,t+1} = \frac{L_{is,t+1}}{L_{ist}}$ and $\dot{L}_{i,t+1} = \frac{\sum_s L_{is,t+1}}{L_{it}}$.
2. Compute industry employment growth by skill group. Start by computing changes in industry choice shares by education group, which can be done by using given contemporaneous wage changes together with past industry employment shares
$$\dot{\pi}_{is\iota t} = \frac{\pi_{is\iota,t-1}\dot{W}_{i\iota t}^{a}}{\sum_{\iota=1}^{N_{ind}} \pi_{is\iota,t}\dot{W}_{i\iota t}^{a}}$$
Then, industry employment changes are given by
$$\dot{L}_{is\iota,t+1} = \dot{\pi}_{is\iota t}\dot{L}_{is,t+1}, \tag{49}$$
and levels can be updated by computing $L_{is\iota,t+1} = \dot{L}_{is\iota,t+1}L_{is\iota,t}$ and $L_{i,\iota,t+1} = \sum_s L_{is\iota,t+1}$.
3. Update local industry-level wages by computing wage growth from agglomeration $\dot{W}_{\iota i,t+1} = (\dot{L}_{i,\iota,t+1})^{\tilde{\eta}_{LD}}$, and then applying it to the previous wage levels: $W_{\iota it} = \dot{W}_{\iota i,t+1}W_{\iota it}$.
4. Aggregate the industry-level wage changes to changes in the overall wage option index for city i by using the fact that
$$\dot{\widetilde{W}}_{ist} = \left(\sum_{\iota} \pi_{\iota ist}\left(\dot{W}_{\iota it}\right)^{a}\right)^{\frac{1}{a}}.$$
5. Update house prices by first computing updated housing expenditure
$$HD_{i,t+1} = \sum_{\iota}\sum_{s} \alpha_s W_{\iota i,t+1}L_{is,\iota,t+1},$$
and housing expenditure growth $\dot{HD}_{i,t+1} = \frac{HD_{i,t+1}}{HD_{it}}$, and then computing house price growth as
$$\dot{P}_{i,t+1} = \left(\dot{HD}_{i,t+1}\right)^{\eta_H^i}$$
6. Now, location utility growth is implicitly defined by
$$\mu_{s,t+1}^{ik} = \frac{\mu_{st}^{ik}(\dot{\mathcal{U}}_{t+1}^{ks})^{\frac{1}{\theta}}}{\sum_j^N \mu_{st}^{ij}(\dot{\mathcal{U}}_{t+1}^{js})^{\frac{1}{\theta}}}$$
$$\dot{\mathcal{U}}_{t+1}^{is} = \dot{A}_{is,t+1}\dot{W}_{is,t+1}(\dot{Q}_{i,t+1})^{-\alpha_s}\left(\sum_k^N \mu_{s,t+1}^{ik}(\dot{\mathcal{U}}_{t+2}^{ks})^{\frac{1}{\theta}}\right)^{\beta\theta},$$

assuming $\dot{Q}_{i,t+1} = \dot{P}_{i,t+1}$, and on the steady-state equilibrium path without endogenous amenities $\dot{A}_{is,t+1} = 1$. Substituting, we can write this as

$$\dot{\mathcal{U}}^{is}_{t+1} = \dot{A}_{is,t+1}\dot{W}_{is,t+1}(\dot{Q}_{i,t+1})^{-\alpha_s}\left(\sum_k^N \left(\frac{\mu^{ik}_{st}(\dot{\mathcal{U}}^{ks}_{t+1})^{\frac{1}{\theta}}}{\sum_j^N \mu^{ij}_{st}(\dot{\mathcal{U}}^{js}_{t+1})^{\frac{1}{\theta}}}\right)(\dot{\mathcal{U}}^{ks}_{t+2})^{\frac{1}{\theta}}\right)^{\beta\theta},$$

which defines a system of equations with N unknowns (the elements of the vector $\dot{\mathcal{U}}^{is}_{t+1}$) and N equations, as $\dot{\mathcal{U}}^{ks}_{t+2}$ is taken as given within each iteration. This can be solved by an equation solver for the elements of $\dot{\mathcal{U}}^{is}_{t+1}$ that correspond to

$$\dot{\mathcal{U}}^{is}_{t+1}\left(\sum_j^N \mu^{ij}_{st}(\dot{\mathcal{U}}^{js}_{t+1})^{\frac{1}{\theta}}\right)^{\beta\theta} - \dot{A}_{is,t+1}\dot{W}_{is,t+1}(\dot{Q}_{i,t+1})^{-\alpha_s}\left(\sum_k^N \mu^{ik}_{st}(\dot{\mathcal{U}}^{ks}_{t+1})^{\frac{1}{\theta}}(\dot{\mathcal{U}}^{ks}_{t+2})^{\frac{1}{\theta}}\right)^{\beta\theta} = 0 \ \ \forall i$$

7. Update $\dot{\mathcal{U}}^{ks}_{t+2}$ based on the newly computed series of $\dot{\mathcal{U}}^{is}_{t+1}$, but always setting it equal to one at time T.

This loop is repeated until the series $\dot{\mathcal{U}}^{is}_{t+1}$ do not change anymore between iterations, which is equivalent to population growth being the same across iterations. The allocation of populations in this economy at time T is then the steady state.

G.3 Imputing historical mobility matrices and amenities

Based on an existing time series of utility growth from some time T_1 forward, a mobility matrix at time T_1 and a full time series of within-period prices and population allocations, we can infer historical mobility matrices using the following algorithm. which solves backwards for $\mu^{ik}_{st}\ \forall i,k$ and $\{\dot{A}_{ist,t+1}\}$, if we're given $\mu^{ik}_{s,t+1}\ \forall i,k$, and $(\dot{\mathcal{U}}^{js}_{t+2})^{\frac{1}{\theta}}$, and the time series of total population, wages and house prices for all past periods.

1. The location choice updating equation implicitly defines a system of equations that relates past migration choices to future choices and future utility:

$$\mu^{ik}_{st} = \frac{\mu^{ik}_{s,t+1}}{(\dot{\mathcal{U}}^{ks}_{t+1})^{\frac{1}{\theta}}}\left(\sum_j^N \mu^{ij}_{st}(\dot{\mathcal{U}}^{js}_{t+1})^{\frac{1}{\theta}}\right) \tag{50}$$

For simplicity, define $\dot{R}^w_{is,t+1} = \dot{\widetilde{W}}_{is,t+1}(\dot{Q}_{i,t+1})^{-\alpha_s}$ as the real wage index. Then, we want to solve the utility growth expression

$$\dot{\mathcal{U}}^{is}_{t+1} = \dot{A}_{is,t+1}\dot{R}^w_{is,t+1}\left(\sum_k^N \mu^{ik}_{s,t+1}(\dot{\mathcal{U}}^{ks}_{t+2})^{\frac{1}{\theta}}\right)^{\beta\theta},$$

jointly with equation 50 for the implied values of $\mu^{ik}_{st}\ \forall i,k$ and $\{\dot{A}_{is,t+1}\}$. These equations implicitly define period t migration flows as functions of unobservable amenity changes. Here, we have $SN\times(N+1)$ unknowns (amenity changes for each location and group, and city-to-city flows for each group), and $SN\times N$ equations, so we need to impose an additional SN conditions to solve for the unknown variables.

To obtain these conditions, I first make use of the population constraints:

$$L_{it} = \sum_s^S L_{ist}$$

$$L_{is,t+1} = \ell_{i,t+1} + \sum_k^N \mu^{ki}_{st}L_{ist},$$

which add the series of SN population levels by group to the unknowns, but provide $(S+1)N$ equations, so we gain N additional constraints. To gain an additional N constraints, we could make additional assumptions regarding the amenity growth process – for instance, assuming $\dot{A}_{ks,t+1} = \dot{A}_{k,t+1}\ \forall k$. However, as we are interested in group differences in preferences for cities, this is not appropriate in this setting.

Instead, I make use of the fact that I have information on *aggregate* migration flows from IRS data for the period where flows by education group are unavailable. Therefore, I can impose

$$\mu_t^{ii} L_{it} = \sum_s^S \mu_{st}^{ii} L_{ist}.$$

That is, I constrain the number of people not moving in each group to add up to the known total number of people staying in place. This yields an additional N constraints and allows me to identify the unknown variables. As a result, these equations can be solved exactly for $\{\mu_{st}^{ik}\}_{i,k\in N}$ and $\{\dot{A}_{i,t+1}\}_{i\in N}$.

2. Using these expressions, I solve for $\mu_{st}^{ik}\ \forall i,k$ and $\{\dot{A}_{ist,t+1}\}$ in the following manner: , choose a candidate vector $\{\dot{A}_{ist,t+1}\}$ for each group s, and substitute into the expression for $\dot{\mathcal{U}}_{t+1}^{is}$. Then, solve for μ_{st}^{ik} conditional on $\dot{\mathcal{U}}_{t+1}^{is}$. Evaluate the additional constraints on population sizes and gross migration rates for this mobility matrix. Additionally, I impose the constraint that all population values have to be positive. Iterate on the initial choice of $\{\dot{A}_{ist,t+1}\}$ until all the constraints are met, using a nonlinear equation solver.
3. Once $\{\mu_{st}^{ik}\}_{i,k\in N}$ and $\{\dot{A}_{it}\}_{i\in N}$. are known, we can compute $\dot{\mathcal{U}}_{t+1}^{is}$ from the equation shown above.
4. This process is then repeated until some period $T0$ is reached.

Note that this imputation of mobility does not make use of any mobility data for periods before $T1$. Therefore, a good check for the model fit is the comparison of partial or aggregate flows at T0 (or intermediate periods) between locations to those implied by the backward simulation of the model.

G.4 Implementing counterfactual migration cost changes

In order to implement the counterfactual changes in migration costs, I start from a rewriting of the logit form for the probability of deciding to relocate to any city k from city i given by Equation 14:

$$\mu_{st}^{ik} = \frac{\exp(\upsilon_t^{ks})^{\frac{1}{\theta_s}}\,\tilde{\tau}_s^{ik}}{\sum_j^N \exp(\upsilon_t^{js})^{\frac{1}{\theta_s}}\,\tilde{\tau}_s^{ij}},$$

where $\tilde{\tau}_s^{ik} = exp(-\tau_s^{ik})^{\frac{1}{\theta_s}}$ is akin to a "migration discount factor" on the utility obtained from other locations. Assume that this factor can be decomposed further into an element related to the personal and psychological cost of leaving the current city of residence for any other location $\tilde{\tau}_s^{ik,\text{Leave}}$, as well as a component related to the cost of changing states $\tilde{\tau}_s^{ik,\text{State}}$, and a bilateral component containing any other bilateral moving costs, including those shown in the estimates in Table 2:

$$\tilde{\tau}_s^{ik} = \tilde{\tau}_s^{ik,\text{Leave}} \cdot \tilde{\tau}_s^{ik,\text{State}} \cdot \tilde{\tau}_s^{ik,\text{Other}}.$$

Moreover, the cost factors of leaving the current city and state are assumed to be the same, no matter which new state or city is chosen, and set to one (i.e. no cost) if there is no change.

Then, the different migration cost scenarios in counterfactual 1 are implemented as follows: The "higher mobility" scenario corresponds to an increase by a factor $\lambda_{cf1} > 1$ in the valuation of locations outside of the current city, i.e. I set

$$\tilde{\tau}_{s,new}^{ik,\text{Leave}} = \lambda_{cf1}\tilde{\tau}_{s,old}^{ik,\text{Leave}}$$

at the beginning of the simulated counterfactual time series. As a result, the baseline migration probabilities change such that

$$\mu_{s,new}^{ik} = \frac{\exp(\upsilon_t^{ks})^{\frac{1}{\theta_s}}\lambda_{cf1}\tilde{\tau}_{s,old}^{ik}}{\exp(\upsilon_t^{is})^{\frac{1}{\theta_s}} + \sum_{j:j\neq i}^N \exp(\upsilon_t^{js})^{\frac{1}{\theta_s}}\lambda_{cf1}\tilde{\tau}_{s,old}^{ij}},\quad \forall\, k \neq i$$

Substituting from the definition of $\mu_{s,old}^{ik}$, this is equivalent to

$$\mu_{s,new}^{ik} = \frac{\lambda_{cf1}\mu_{s,old}^{ik}}{\mu_{s,old}^{ii} + \sum_{j:j\neq i}^N \lambda_{cf1}\mu_{s,old}^{ik}}\quad \forall\, k \neq i.$$

That is, this change in the fundamental migration cost is equivalent to scaling all observed migration probabilities to other cities by the same factor and then rescaling to ensure migration probabilities add to one.

For the "no inter-state migration" scenario, the implementation corresponds to setting the between-state migration cost component to infinity, i.e. the migration discount factor goes to zero:

$$\tilde{\tau}_s^{ik,\text{State}} = 0 \Leftrightarrow state_i \neq state_k.$$

As a result, it will be the case that

$$\mu_{s,new}^{ik} = 0 \Leftrightarrow state_i \neq state_k.$$

For the "no migration scenario" it is then the case that

$$\tilde{\tau}_{s,new}^{ik,\text{Leave}} = 0 \Leftrightarrow i \neq k,$$

such that

$$\mu_{s,new}^{ik} = 0 \Leftrightarrow i \neq k.$$

In Section 8.6, I consider the following scenarios: (A) The cost of leaving the current city of residence falls by half ($\tilde{\tau}_s^{ik,\text{Leave}}$ doubles if $i \neq k$). (B) Inter-state migration costs become prohibitively large, i.e. $\tilde{\tau}_s^{ik,\text{State}} = 0$ if i and k are in different states. (C) Migration costs to any other city become prohibitively large ($\tilde{\tau}_s^{ik,\text{Leave}} = 0$ if $i \neq k$).

Because of the direct mapping between migration costs and baseline migration probabilities, these scenarios are implemented in my analysis as follows: (A) In the increased mobility scenario, I multiply all migration probabilities μ_{st}^{ik} that represent leaving the origin city (so $i \neq k$) by 1.5 in order to represent a shock to moving costs that results in an average increase of gross migration by 50%. (B) To implement the no-interstate migration scenario, I set all migration shares μ_{st}^{ik} to zero if i and k are not in the same state. This is equivalent to distributing the migration shares directed to out-of-state destinations between all in-state cities according to their relative share of original migration flows. (C) To simulate no migration, I simply set $\mu_{st}^{ii} = 1 \ \forall i$. In all scenarios, I reweight all moving probabilities after the adjustment, so that they add to one.

H Supplementary Analyses

H.1 Descriptive analysis of drivers of city attractiveness over time

While persistent migration costs determine *which* cities are more likely to be part of the same migration network, the variation over time in the level of migrant flows into or out of a city will depend on changes in the overall attractiveness of that city. In order to get a descriptive sense of what factors correlate with migration flows, and to supplement the causal effect estimates in Section 4, this section details non-causal estimates of the drivers of bilateral flows between cities.

As a first pass at determining which city characteristics have been associated with more or less migration, I take the destination city × year fixed effects obtained from the migration gravity regression shown in equation 1 and regress them on city characteristics, a national trend α_t, and – in some specifications – a city fixed effect α_j:

$$\theta_{jt} = \alpha_t + \alpha_j + \alpha_w \ln W_{jt} + \alpha_P \ln P_{jt} + \beta' \mathbf{X_j} + \epsilon_{ijt}$$

The included time-varying characteristics are the average local wage income (from the IRS), a quality-adjusted house price index,[115] and indices of local consumption amenities (constructed as principal components of the density of establishments in leisure and consumption industries).[116] In addition, I also include time-invariant measures of the *level* of the cultural and natural amenities discussed in the previous section, and population density.

Column 1 of Table A7 shows the effect of these characteristics on the city attractiveness series obtained from the 1990-2017 aggregate migration sample. Column 2 adds a city fixed effect, which absorbs the effect of any unobserved time-invariant city characteristics. The results show that greater overall migration to a city destination

[115] I adjust the level of house prices for the quality of the housing stock in the year 2000, following Albouy and Ehrlich (2018), and then use the FHFA repeat-sales index to adjust house prices for changes over time. See Appendix F.7 for details.

[116] See Appendix F.4 for details

Table A7: City attractiveness determinants

Sample:	1990-2017		2005-2017			
	All	All	Coll.	Coll.	Non-Coll.	Non-Coll.
	(1)	(2)	(3)	(4)	(5)	(6)
Log Inc. per cap.	2.212***	0.701***	4.923***	1.714***	4.541***	1.739***
	(0.298)	(0.095)	(0.486)	(0.299)	(0.478)	(0.312)
House price index (qual.-adj.)	0.699***	-0.202***	0.677***	-1.142***	0.653***	-1.215***
	(0.121)	(0.050)	(0.211)	(0.156)	(0.205)	(0.157)
Amenities index (1st PC)	-0.378***	-0.011**	-0.399**	0.030	-0.379**	0.030
	(0.088)	(0.005)	(0.164)	(0.022)	(0.158)	(0.021)
Amenities index (2nd PC)	0.291***	-0.030***	0.441***	-0.072*	0.444***	-0.079**
	(0.048)	(0.010)	(0.085)	(0.039)	(0.084)	(0.039)
Water surface share	0.762*		0.786		0.831	
	(0.390)		(0.655)		(0.634)	
Nontrad. Christ. Share	0.162		0.416		0.431*	
	(0.168)		(0.265)		(0.256)	
Jan. Temperature	0.028***		0.035***		0.036***	
	(0.003)		(0.006)		(0.006)	
College share (2000)	3.406*		0.949		-5.890*	
	(1.793)		(3.169)		(3.061)	
Population density	0.436*		0.789**		0.565	
	(0.241)		(0.399)		(0.370)	
Observations	19,208	19,208	9,308	9,308	9,308	9,308
Year FE	✓	✓	✓	✓	✓	✓
City FE		✓		✓		✓

Heteroskedasticity-robust standard errors clustered at the CZ level in parentheses: * p<0.10, ** p<0.05, *** p<0.01. The table shows results from regressions where the dependent variable consists of CZ × year fixed effects from a migration gravity regression. Analysis includes all continental U.S. CZs, excl. New Orleans, for which fixed effects could be computed, leading to a total of 613 - 716 CZs in a given year. See text for description of the explanatory variables.

is significantly correlated with higher average wages in the city. Moreover, I also find significantly positive effects of a warmer climate, and there is weak evidence that greater density, college shares and water access are associated with greater aggregate migration. Once constant unobservable city attractiveness is accounted for, house prices show the expected negative sign in Column 2. Moreover, amenities in the form of consumption establishments have a significant correlation with migration in both specifications.

These descriptive findings inform the empirical approach in later sections. In particular, this analysis suggests that one needs to be wary of bias arising from time-varying changes in amenities, when estimating the effect of wages and house prices on location choices.

Heterogeneity in correlates of city attractiveness. In columns 3 to 6, I analyse the determinants of city attractiveness separately for the city-year fixed effects obtained from migration flows by education group. The estimates in columns 3 and 5 show that the main differences between the two education groups are that only college-educated workers are significantly more likely to migrate into denser cities, while the college share has a weak negative effect on migration of non-college workers. In this descriptive analysis, I only find small differences in wage and house price correlations with migration by education group. However, comparing the aggregate migration coefficients and the separate estimates by education group shows that composition bias might lead to very different estimates when taking into account the subgroup heterogeneity.[117] The quantitative model in a later section will account for this subgroup heterogeneity, as well as the effect of observed and unobserved changes in amenities when trying to estimate the location choice parameters.

H.2 Migration links as predictors of inter-city house price correlation: horserace regressions

In Section 4.3.2, I built on the analysis in Sinai and Souleles (2013) and showed that migration links perform well as predictors for bilateral correlations in house price growth between cities when compared to other measures individually. In this section, I provide further details on the alternative link measures and show that migration links are strong predictors of house price correlation even when controlling for *all* the alternative links jointly in horserace regressions.

To put the ability of migration links to predict house price correlations in context, I consider a number of alternative measures of inter-city links: First, I include an inverse-distance weighted measure that represents the notion that house price correlations might stem from shocks that are common among cities that are geographically close to one another.

Second, I consider a social connectedness index (SCI) based on Facebook friendship links between geographic areas that was introduced in Bailey et al. (2018b). In a related paper, Bailey et al. (2018a) showed that differences in individual exposure through online social networks to house price movements in distant counties can predict differences in housing investment decisions. To measure the importance of this alternative channel, I use weights based on the SCI measure to construct a social network-weighted measure of house price correlations.[118]

Third, I construct a destination population-weighted measure of house price correlations to control for the possibility that the migration weights are simply picking up the fact that large cities have more migration links and might be driving the housing cycle of smaller cities.

Fourth, I explore the possibility that house price correlations reflect similarity in industry structure between cities. To measure industry structure differences, I compute the vector distance in 2-digit NAICS industry employment shares by city, and use the inverse of this distance to measure similarity in industry structures.

Fifth, I include an equal-weighted measure that simply reflects a city's average correlation with other cities' house prices.

We can test the predictive ability of these different city links more formally by estimating the regression model

$$corr(\Delta \ln P_i, \Delta \ln P_k) = \alpha + \beta_1 MigShare_{i \to k} + \beta_2 \ln dist_{ik} + \beta_3 SCI_{ik} + \beta_4 IndDist_{ik} \theta_i + \theta_k + \epsilon_{ik},$$

[117]Note that the full sample estimates also differ in the sample period, in addition to varying the level of aggregation. However, in regressions not reported here, I find that even when restricting the aggregate sample to 2005-2017, the estimates differ from the subgroup estimates.

[118]The SCI index measures the relative probabilities of friendship links between counties (normalized for their respective Facebook user base) – which I aggregate to the commuting zone level, and then normalize for each CZ such that the weights for all other CZs sum to one.

which represents a horserace between migration links, distance, social networks, and industry structure in explaining variation in house price correlations across CZs.

Table A8: City links and house price growth correlation

Dependent variable:	House price growth correlation coeff. × 100			
	(1)	(2)	(3)	(4)
Panel A: Full migration network				
Migration outflow share	83.276***	69.339***	8.165**	10.378***
	(5.415)	(4.372)	(3.597)	(3.568)
Log distance (miles)			-5.846***	-5.405***
			(0.216)	(0.211)
Social Connectness Index			10.327	10.145
			(7.347)	(7.268)
Industry structure similarity				22.573***
				(2.446)
Observations	516,241	516,241	516,241	516,241
R-squared	0.00	0.65	0.67	0.68
Panel B: Distance > 150 **mi.**				
Migration outflow share	235.132***	150.743***	72.698**	66.976**
	(80.653)	(50.749)	(29.852)	(27.499)
Log distance (miles)			-4.812***	-4.310***
			(0.242)	(0.243)
Social Connectness Index			611.968***	602.176***
			(69.872)	(70.143)
Industry structure similarity				22.413***
				(2.446)
Observations	495,554	495,554	495,554	495,554
R-squared	0.00	0.66	0.67	0.68
Origin FE		✓	✓	✓
Destination FE		✓	✓	✓

Heteroskedasticity-robust standard errors two-way clustered at the origin and destination CZ level in parentheses: * p<0.10, ** p<0.05, *** p<0.01. Correlations in house prices growth and average outmigration shares are calculated over 1990-2017. Data contains pairs of all 721 continental U.S. CZs, excl. New Orleans, that have any known outmigration flows.

The results are shown in Appendix Table A8. Going from left to right, the columns add in alternative measures of inter-city links as well as origin and destination city fixed effects that capture the overall tendency of cities to be more connected (e.g. due to city size). As the table shows, migration flows have a significant positive association with the correlation in house prices. Moreover, this relationship is robust to controlling for spatial correlation due to geographic proximity as well as variation associated with social network links and industry similarity. Panel A shows the results for the full migration network whereas Panel B limits the sample to those city pairs which are at least 150 miles apart, which corresponds to the long-distance migration network used in my baseline regressions.

Even in the most stringent specifications in column 4, migration links continue to be significant predictors of house price correlations. That is, migration links contain information about city house price connections that go beyond the set of links represented by the other links included in the regressions.

The estimated (non-causal) coefficient in Column 4 for the long-distance network indicates that a 10 ppt greater share of migrants from city i going to city k is associated with a 6.7 ppt higher correlation in house price growth between the cities, holding constant the city's co-movement with house prices overall. Incidentally, this is similar to the long-run house price spillover estimates found in the baseline reduced form regressions.

Trade flows and migration flows. In the graphs showing the individual power of different measures of links in predicting pair-wise house price correlations, I also show results using a measure of trade flow links between cities that represents industry linkages and the propagation of economic shocks through input-output networks (see Appendix Section F.9 for details on how this measure is constructed). However, this measure is only available for a limited subsample corresponding to less than 10% of the pairwise links and is therefore not included in

Figure A16: Trade flows and migration flows by location pair. Panel (a) shows the value of 2012 log gross migration flows (sum of flows in both directions) over the log gross value of 2012 trade flows (in both directions) for each pair of CFS areas. Panel (b) shows the same variables but residualized with regard to the log sum of populations of the location pair in the year 2000 as a proxy for city sizes. The size of location pair markers is shown proportional to the joint population size of the pair, and the graphs also show the unweighted line of best fit.

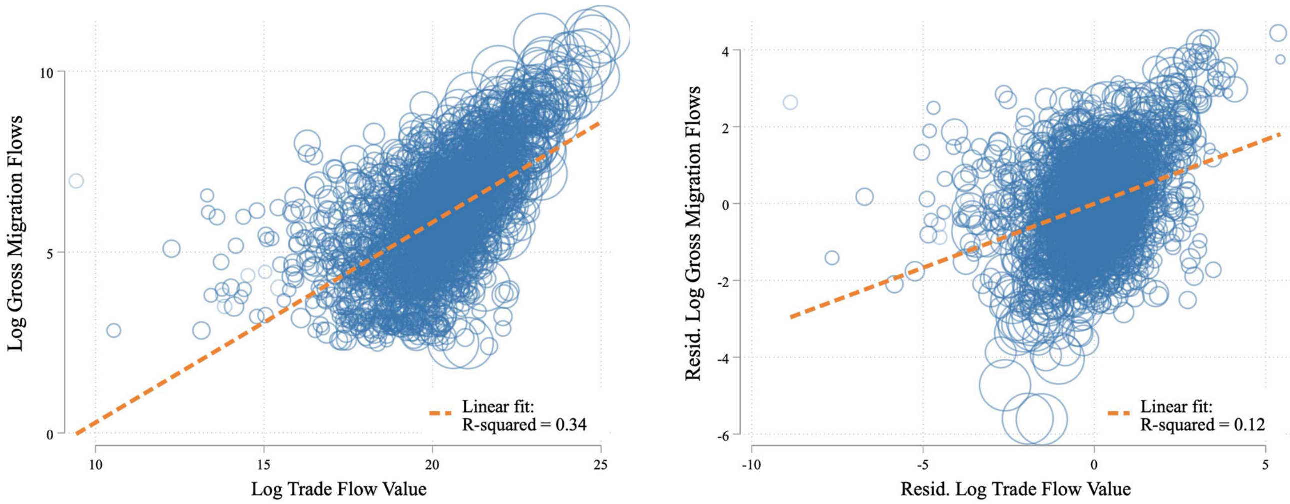

(a) Raw flows

(b) Resid. w.r.t. population size

the regressions. In unreported regressions, I find that, on that limited sample, migration links are significant in predicting house price correlations when controlling for trade flow links.

In order to further explore the degree to which migration links overlap with trade flows, consider Appendix Figure A16, where I show how gross migration flows and the value of trade flows are correlated for pairs of Commodity Flow Statistics areas (which are on average somewhat larger than CZs). This data is available for all connections between 70 CFS areas, (although the graphs drop any pairs of locations that have no migration flows or no trade flows).

As the graphs show, while there is a relatively high R-squared of 34% between migration flows and the value of trade flows between two locations in the raw data, this correlation is to a large degree driven by the fact that large cities are more likely to trade and have migration flows with *any* other city. Once I adjust for the joint population size in each location pair, the R-squared is only 12%. In conjunction with the fact that – as shown in Figure 10 – migration links are a better predictor of house price correlations than trade flows, this makes it unlikely that migration links are merely a proxy for trade links that are truly driving the observed house price spillover effects.

H.3 Reasons for moving and demographics of movers

In order to understand better which demographics are driving migration patterns in the U.S., this section will use data from the Current Population Survey (CPS) to explore the characteristics of U.S. migrants.

Reasons for moving. First, I consider the reasons stated by survey respondents who moved in the last year when asked why they were moving. I group the CPS response categories into "family reasons" (*change in marital status; establishing own household; other family reasons*), "employment reasons" (*new job or job transfer; to look for work or lost job; to be closer to work/easier commute; other job-related reason*), "retirement", and "housing reasons" (*wanted to own home, not rent; wanted new or better house/ apartment; wanted better neighborhood/less crime; wanted cheaper housing; foreclosure/eviction; other housing reason*). I omit move reasons due to college attendance, climate, health, or disaster, which the CPS groups under "other reasons". For each category of reasons, I compute the share of the total population with available moving status that moved for that reason - distinguishing between all moves and moves across county boundaries.

The patterns for moving reasons are shown in Appendix Figure A17, in Panels (a) and (b). As the time series show, housing reasons are the most important reason for moving by far when I include within-county moves in the tabulation in Panel (a), followed by family reasons and employment reasons. When I consider only "long-distance" moves that cross county boundaries, housing, employment and family reasons are of similar importance. In either category of moves, moves due to retirement play a negligible role in mobility. The time pattern of migration being correlated with housing booms and busts is reflected in the fact that moves for any reason are higher in the early 2000s and 2010s than in the late 2000s. However, it is important to note that this is mostly driven by the cyclicality of moves for housing reasons. This provides support for the idea that migration cycles are an important consequence and component of housing cycles rather than being driven by changes in employment opportunities that coincide with housing booms.

It is important to note the limitations of this analysis: in a spatial equilibrium setting with Rosen-Roback style preferences, like the model in this paper, residents jointly consider the effect of employment, housing, family and other amenities on their utility in deciding whether or not to change locations – so there is no clear sense in which either of these elements is *the* cause of their move. However, if we assume that respondents take the survey to be asking about which of these elements had been changing the most to occasion their change in location preference, then the prominent role of housing provides qualitative evidence of the proposed migration spillover mechanism where house price changes lead to migration in search of more affordable housing during housing booms.

Employment status of long-distance movers. Next, I consider the employment status of movers. In line with the paper's focus on moves across cities, I focus on inter-county movers. I retain all CPS respondents who moved across county lines and have an employment status (which drops children, for instance), and are not in the armed forces. Then, I compute the number of movers in 3 non-overlapping and exhaustive categories: (1) Employed (both at work currently and not currently at work); (2) Unemployed, or not in the labor force ("NILF"), but not retired; (3) NILF and retired. The share of inter-county movers in each employment category is shown in Panel (c) of Appendix Figure A17. Around 60% of all movers are employed (when being surveyed after their move), less than 10% are retired, and the remainder are unemployed or out of the labor force but not retired. This aggregate pattern is particularly important when evaluating anecdotal evidence about particular city pairs with strong migration links. For example, it is possible that a greater number of movers from New York City to Florida are retirees than for other city pairs, but given that retirees represent a very small share of migration overall, it is unlikely that they represent the majority share of moves even for this city pair. More generally, it is unlikely that retirement location preferences play a major role in explaining the migration spillover patterns documented in this paper.

Age structure of long-distance movers. I next consider the age structure of long-distance movers, grouping CPS respondents by age, and omitting those less than 20 years old. The age shares of movers are shown in Panel (d) of Appendix Figure A17. The time series show that mobility declines precipitously with age: while the population share of those aged 40 years and older is larger than that of those under 40, their combined share of long-distance migration is less than half that of the younger group. While there is a gentle upward slope in the migration share of older groups, this can likely be explained by their increasing population share over this time period. As a result, when we are thinking about migration patterns, they are likely to be driven predominantly by the decisions of the younger working-age population.

H.4 Stationary steady-state equilibrium example: year 2005 data

As examples of what the stationary steady state equilibria look like, I compare the underlying observed data for 2005 to the eventual SSE starting at 2005.

I first map the baseline of the *actual* college share in 2005 and house price growth 1990-2005 in each CZ in the continental U.S, which is shown in Panels A and B of Figure A18. The maps show that college shares were particularly high in 2005 in metropolitan areas across the U.S., and in the coastal areas of New England, Florida, and California, while house price growth up to the peak of the boom was concentrated in the Northwest, Florida, and the Atlantic Coast.

Then, to illustrate the steady state concept of the economy "settling" into equilibrium based on the observed characteristics in a given year, I map the changes between the observables for 2005 and the steady state values for 2005 in house prices and college shares in Panels C and D of Figure A18. The maps show that on the path to the steady state equilibrium consistent with 2005 observables, college shares are predicted to increase further in many areas where they were already high. In addition, college shares are predicted to increase in some low-college share areas, in particular Nevada and Arizona, as well as the Pacific Northwest. These predicted changes between 2005 actuals and the 2005-implied steady state can be interpreted as the equilibrium that the 2005 disequilibrium values were moving towards and would reach if no other shocks occurred. It's important to distinguish this from

Figure A17: Reason for moving and demographics of movers (CPS). The data shown in the graphs below comes from the Current Population Survey March Supplement. Panel (a) shows the share of the U.S. population moving for the stated reason, omitting the category of "other reasons". Panel (b) considers only "long-distance" moves that cross county boundaries. Panel (c) shows the share of inter-county movers by employment status, distinguishing between the employed, retirees and the remainder of workers who are unemployed or not in the labor force. The analysis excludes workers in the armed force or where employment status is not reported. Panel (d) groups respondents by their age and plots shares of total inter-county migration for each age group.

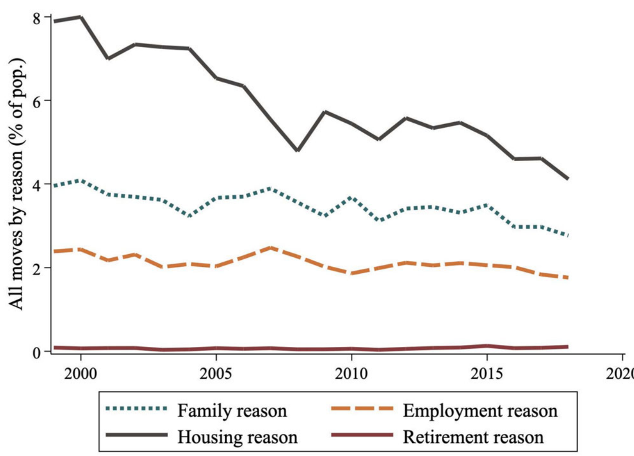

(a) Reason for moving: all movers

(b) Reason for moving: inter-county movers

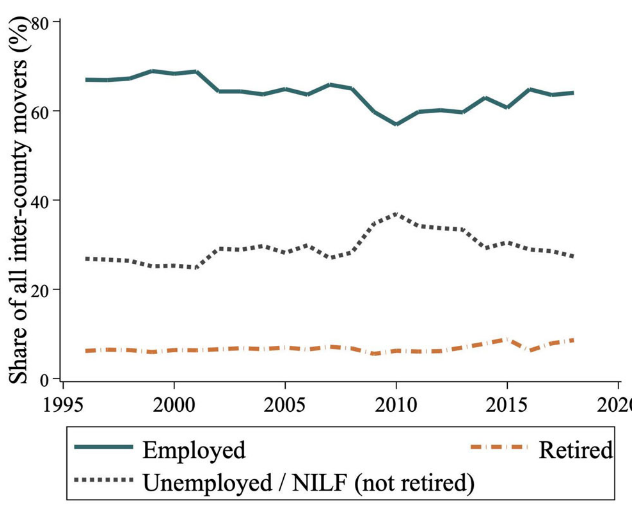

(c) Inter-county mig. by employment status

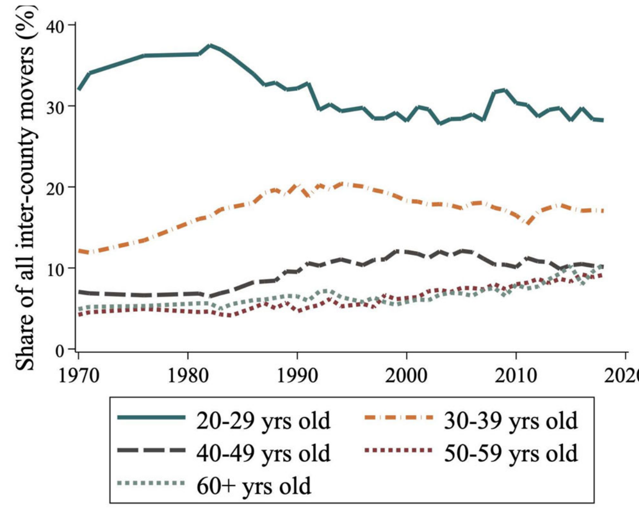

(d) Inter-county mig. by age

Figure A18: Stationary steady-state equilibrium example: 2005. The map in Panel A and Panel B plot actual FHFA repeat-sales house price index growth 1990-2005 and college population shares in 2005 as the historical observables before the transition to the stationary steady state is computed. The maps in Panels C and D then plot the growth rate in college shares and house prices in the model-predicted steady state equilibrium consistent with 2005 observables relative to the actual values in 2005. All maps display values for 512 Adjusted CZs overlayed on 721 continental U.S. CZs, omitting New Orleans.

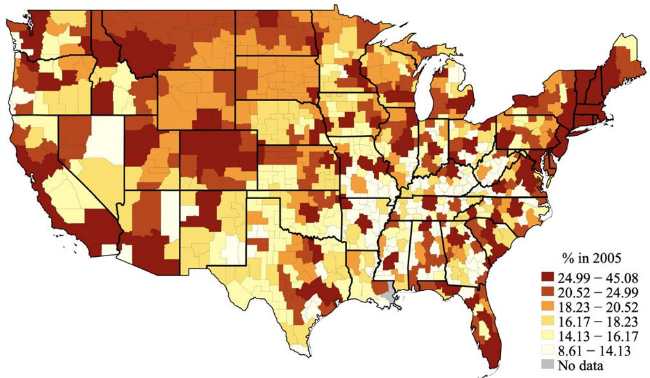

(a) College shares in 2005

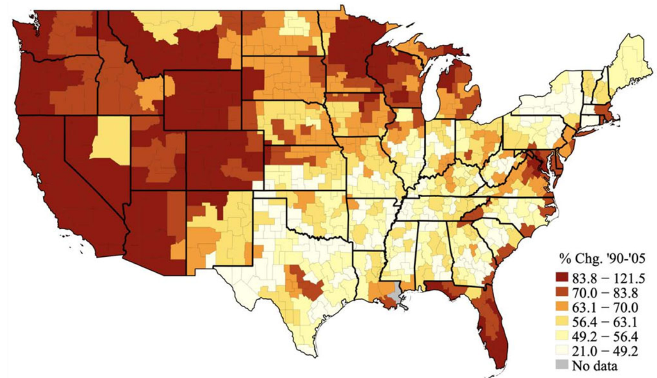

(b) House price growth 1990-2005

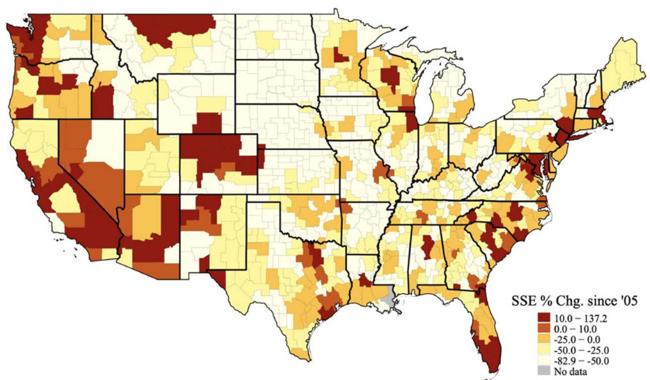

(c) College share growth: 2005 SSE vs. 2005 Actuals

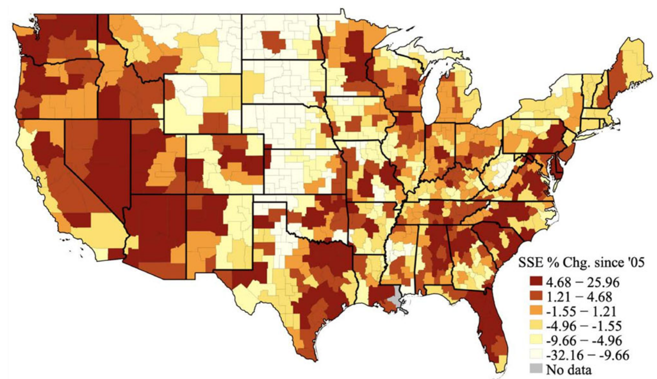

(d) House price growth: 2005 SSE vs. 2005 Actuals

the predicted path of the actual economy: by design the steady-state equilibrium for 2005 does not account for the fact that a financial crisis and the Great Recession were about to hit the U.S. economy a few years later.

H.5 Superstar cities and displacement during the 2000-2007 boom

In this section, I explore how the shifts in educational composition accompanying housing booms that are documented in Section 4.2.3 played out during the 2000-2007 housing boom for superstars and migration spillover cities in particular. First, the relationship between population growth and changes in the college share is plotted for all large CZs in the left panel of Appendix Figure A19. While all large CZs saw an increase in their college share during this period,[119] this compositional shift was more pronounced among cities with low population growth. That is, cities that were growing slowly overall, saw a more rapid displacement of non-college workers by college-educated workers. That much of this displacement is associated with crowding out due to housing supply constraints is suggested by the fact that most of the supply-constrained superstar cities experienced high increases in their college share.

[119]This is in part driven by the fact that the share of the college-educated population in the U.S. has continued to grow nationally over the last two decades, albeit more slowly than in previous decades (Autor et al., 2020). Another factor is the greater preference of college-educated workers for large cities relative to non-college educated workers (Couture and Handbury, 2017).

Figure A19: Displacement for superstars. The graph plots average population growth and the change in the college-educated population share over 2000-2007 in the left panel, and changes in the college share and the ratio of house price growth to population growth for the same period in the right panel. Data shown contains 45 CZs on the left, consisting of all continental U.S. CZs, excl. New Orleans, with year 2000 Census population $> 0.85M$, and corresponding to $\sim 50\%$ of the continental U.S. adult population. Right panel additionally excludes 2 CZs with negative population growth from the graph for better visibility.

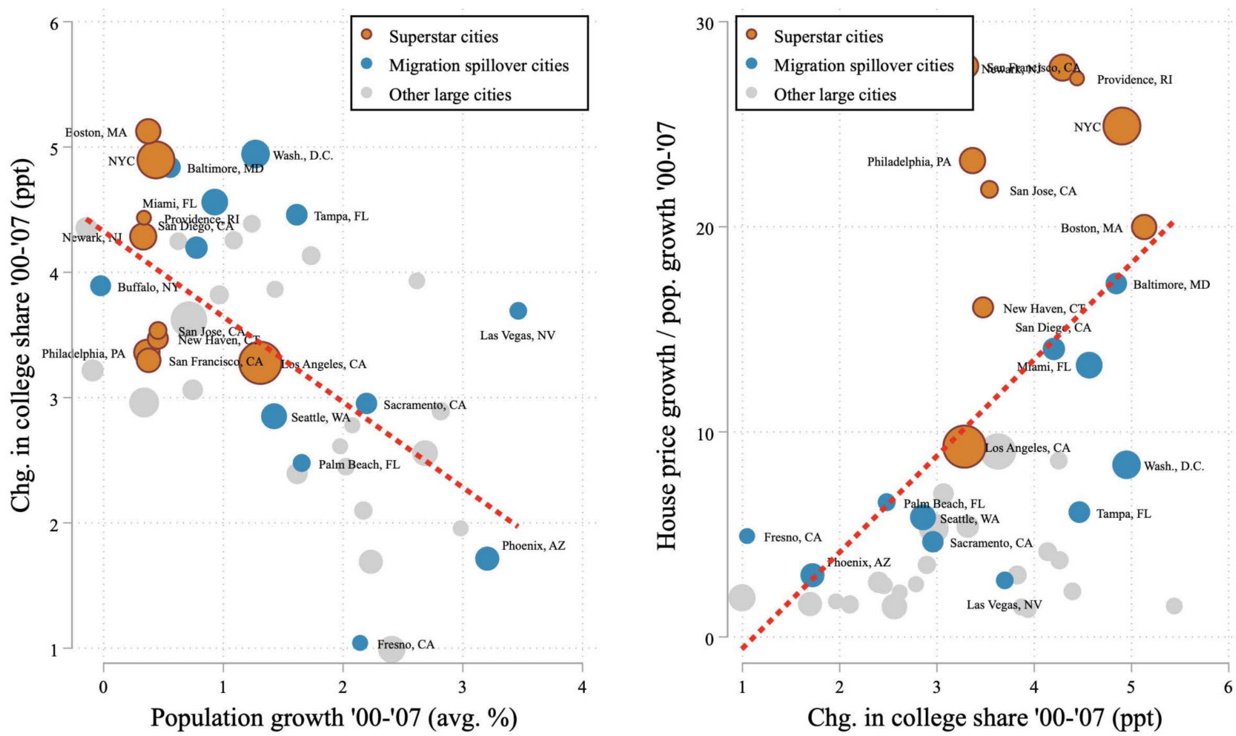

Moreover, most of the migration spillover cities are well above the line of best fit for composition changes. That is, their college shares increased faster than would be expected, relative to other large cities, given their overall population growth.

Next, we can see whether the shift towards a highly educated – and on average higher-income – population can explain the high house price growth seen in superstar and migration spillover cities. In the right panel of Appendix Figure A19, I plot the ratio of house price growth to population growth – as a measure of the slope of the house price response to demand shocks – over the change in the local college share.

The graph shows a strong positive relationship between the degree to which house prices respond to population growth, and the composition of that population growth: if the increase in population involves a greater displacement of low-income residents by high-income residents (here proxied by education levels), then house price growth increases more. In fact, even though we saw before that most of the migration spillover cities experience high levels of house price growth relative to their population growth during the housing boom, this is more than explained by the accompanying shift in demographic composition.

These analyses suggest a mechanism where demographic changes can complement the house price effects of migration spillovers. For instance, focusing again on Las Vegas and Phoenix as salient examples, the left panel shows that they saw relatively large shifts in education levels for their level of population growth. The right panel confirms that the degree to which house prices responded to population growth in these two cities was, if anything, *less* than might be expected once we take into account their rapid demographic shift towards a more college-educated population.