Effects of Immigration on the US Housing Market

Manika Dhiman

Abstract

This paper estimates how immigration affected housing markets in U.S. commuting zones from 2005 to 2019 and whether the effect differs by immigrants' education. Using shift–share instruments built from 1980, 1990, and 2000 settlement shares, I find that an immigrant inflow equal to one percent of the working-age population raises Zillow home values by 3.3–3.7 percent, Zillow rents by 1.4–3.6 percent, and the FHFA house price index by about 1.5 percent, while Fair Market Rents do not respond. The effects come mainly from immigrants without college education, although joint estimates cannot reliably distinguish the two groups.

1 Introduction

Immigration and housing affordability are two of the most contested policy issues in the United States, and they are closely linked. The Congressional Budget Office projects that deaths will exceed births beginning in 2030, after which net immigration will account for all U.S. population growth (Congressional Budget Office, 2026). Population-driven growth in housing demand will therefore come increasingly from immigrants. Because new immigrants need housing as soon as they arrive and settle in a relatively small number of places, this demand is concentrated in particular local markets. At the same time, immigrants work disproportionately in construction, so the same inflows that raise housing demand can also lower building costs and expand housing supply (Monras, 2020). Because these demand and supply effects work in opposite directions, the net effect of immigration on housing costs is theoretically ambiguous. This raises two empirical questions: how does immigration affect local rents and home prices, and does the effect depend on which immigrants arrive?

This paper estimates the effect of changes in the immigrant population on housing markets in U.S. commuting zones between 2005 and 2019, and examines heterogeneity in this effect by immigrants' educational attainment. I use education as a proxy for skill because it shapes income, access to mortgage credit, and the kind of housing people demand. I measure housing outcomes with four series that capture different parts of the market: HUD Fair Market Rents (FMR), which track rents near the 40th percentile of the local rent distribution; the FHFA House Price Index (HPI), a repeat-sales index of constant-quality house prices; the Zillow Observed Rent Index (ZORI), which tracks typical asking rents; and the Zillow Home Value Index (ZHVI), which tracks typical home values.

The main empirical challenge is endogeneity in immigrants' choice of location. Immigrants move to places with growing economies, and those places would see rising rents even without new immigrants, so an OLS regression of housing growth on immigrant inflows would be biased: it would confuse the effect of immigration with the effect of local economic growth. To address this, I use a shift–share instrument. The instrument rests on a well-documented regularity: new immigrants tend to settle where earlier immigrants from the same country settled (Altonji and Card, 1991; Card, 2001). I predict each commuting zone's immigrant inflow by combining its historical settlement pattern by country of origin, measured in 1980, 1990, or 2000, with the national growth of each origin group. Following Goldsmith-Pinkham et al. (2020), identification comes from the historical shares. To assess whether these shares are valid, I test whether they predict housing growth before the study period, control for baseline characteristics that are correlated with them, and use Rotemberg weights to identify the origins that drive the estimates.

I find that immigration raises market rents and home values. With baseline controls, an increase in the immigrant population equal to one percent of the local working-age population raises Zillow home values by 3.3 to 3.7 percent and Zillow rents by 1.4 to 3.6

percent within the same year, depending on the base year of the instrument. The FHFA index rises by about 1.5 percent, while Fair Market Rents do not respond detectably. The Zillow estimates are larger than the estimate of about one percent in Saiz (2007). My FMR estimates are closer in size to Saiz's FMR-based estimate, but they are not statistically significant. Section 6.5 discusses why the estimates differ.

Estimated one group at a time, the housing response comes mainly from immigrants without college education. Their inflows raise Zillow rents by 4.4 to 7.0 percent and Zillow home values by 4.9 to 5.5 percent, and their first stage is adequate in every base year. For college-educated immigrants, the instrument is weak with 1980 and 1990 shares; with 2000 shares, where it is strong, the estimates are small relative to their standard errors. Because the two education instruments are built from the same historical shares, I also estimate both groups jointly. In the joint model with baseline controls, non-college inflows raise Zillow rents by 5.3 to 6.2 percent in every base year, but the joint first stage is weak, and in most specifications I cannot reject that the two groups have equal effects. For all immigrants, interaction models show that the response of Zillow home values rises with the restrictiveness of local land-use regulation in every base year, consistent with demand pressure turning into higher prices where supply cannot expand; the evidence for rents and the FHFA index is weaker. Finally, OLS estimates suggest that the local working-age native population falls by about one person for every three to five working-age immigrants who arrive, but the corresponding 2SLS estimates are imprecise.

This paper contributes to three literatures. First, it adds recent, national evidence to the literature on immigration and local housing markets. Applying the natural-experiment design of Card (1990), Saiz (2003) shows that the 1980 Mariel boatlift, which added about nine percent to Miami's renter population in one year, raised Miami rents by 8 to 11 percent relative to comparison cities, with no effect on higher-quality units. Saiz (2007) finds that an immigrant inflow equal to one percent of a metropolitan area's population raises rents and housing values by about one percent, an effect much larger than the wage effects found in labor markets. Ottaviano and Peri (2006) find that U.S. cities whose foreign-born share rose between 1970 and 1990 saw higher native wages and higher rents, and Mussa et al. (2017) find that immigrant inflows into a metropolitan area are associated with higher rents and house prices both there and in neighboring areas. Outside the United States, immigration accounts for about a third of Spain's 1998–2008 housing boom (Gonzalez and Ortega, 2013), an inflow equal to one percent of a Swiss region's population raises single-family house prices by about 2.7 percent (Degen and Fischer, 2017), and Haifa's housing market absorbed a large and sudden immigration shock in the 1990s with prices adjusting quickly toward their fundamental values (van der Vlist et al., 2011). Within metropolitan areas, by contrast, housing values grow more slowly in the neighborhoods where immigrants settle, a pattern attributed to native flight (Accetturo et al., 2014; Sá, 2015; Saiz and Wachter, 2011). I study a more recent period, cover the whole mainland United States at the commuting-zone level, and estimate the same design on four housing

measures. This comparison shows that the measured effect depends on which part of the market a series measures and where it is observed.

Second, I show how the housing response varies with immigrants' education. Most studies of immigration and housing treat immigrants as one group, although the labor literature has long treated immigrants with different education as different labor inputs (Ottaviano and Peri, 2012). The closest study is Ang et al. (2026), who find that rising Asian population shares raise county house prices and that roughly a third of this effect works through improved school quality. Rather than focusing on one origin group, I divide all immigrants by education. Estimated one group at a time, the short-run housing response comes from immigrants without college education. Estimated jointly, the rent response to non-college inflows remains, but the data cannot distinguish the two groups' effects once baseline characteristics are controlled for.

Third, I apply recent results on shift–share designs (Adão et al., 2019; Borusyak et al., 2022; Goldsmith-Pinkham et al., 2020) to the immigration–housing setting. Sharpe (2019) argues that the standard shift–share instrument overstates the effect of immigration on rents because it does not fully account for immigrants' attraction to growing cities, and shows that controlling for initial city characteristics makes the estimated effect smaller. I follow this approach by interacting baseline characteristics with year effects. Rather than choosing a single base year, I compare instruments built from 1980, 1990, and 2000 shares, test each for pre-trends, and decompose the estimates with Rotemberg weights. The comparison shows that the choice of base year changes the strength of the first stage, the evidence on pre-trends, and which historical settlement patterns drive identification.

The rest of the paper is organized as follows. Section 2 describes how immigration can affect local housing markets. Section 3 describes the data. Section 4 presents the empirical strategy, and Section 5 evaluates the instrument. Section 6 presents the results, Section 7 discusses them, and Section 8 concludes.

2 How Immigration Affects Local Housing Markets

Immigration affects housing markets mainly through demand. New immigrants need housing as soon as they arrive, and most of them rent at first. Immigrant households have lower homeownership rates than native households (Borjas, 2002), so an immigrant inflow raises rental demand first. Over time, some immigrants buy homes, and expected rent growth is capitalized into home values. Saiz (2007) summarizes this sequence: rents rise in the short run, and housing prices catch up.

How much of this new demand shows up in prices depends on how quickly housing supply responds. Where land is scarce or regulation slows construction, supply is inelastic, and a demand shock mostly raises rents and prices (Saiz, 2010). Where supply is elastic, new construction absorbs much of the shock. Immigration can also affect supply directly,

because immigrants work disproportionately in construction. Monras (2020) finds that Mexican immigration after the 1995 peso crisis raised relative rents in the short run, but that in the long run housing prices fell in high-immigration locations, as immigrant workers entered the construction sector and lowered construction costs.

Natives may also respond. If natives leave as immigrants arrive, total housing demand rises by less than the immigrant inflow, and the estimated price effect is net of this offset (Saiz, 2007). Native departures can even lower prices where immigrants settle: housing values grow more slowly in immigrant neighborhoods within U.S. metropolitan areas (Saiz and Wachter, 2011), and immigration lowers house prices across local areas in the United Kingdom as higher-income natives move out (Sá, 2015). Section 6.9 examines how the working-age native population responds.

I hypothesize that these channels differ by education. College-educated immigrants earn more, have better access to mortgage credit, and may become homeowners sooner, while immigrants without college education are more likely to rent for longer and to compete for lower-priced units. In the short run, however, both groups mostly rent, so both can raise rents, especially where supply is inelastic. Over longer horizons, differences in income, credit access, and housing preferences may lead the two groups to affect rents and home prices differently.

Because I estimate the model in annual changes, the estimates in this paper capture short-run responses. Each estimate is the net result of all these channels within a year: the added demand, any construction response, changes in vacancies and household crowding, moves between renting and owning, and native mobility.

3 Data

3.1 Immigration Data

I measure immigration with individual-level microdata from IPUMS USA (Ruggles et al., 2025). For the historical settlement patterns, I use the 1980, 1990, and 2000 decennial census samples. For the study period, I use the annual American Community Survey (ACS) from 2005 to 2019. Both sources record each person's place of birth, citizenship, age, and educational attainment, among other attributes.[1]

I define immigrants as people born outside the United States who were not U.S. citizens at birth, that is, naturalized citizens and non-citizens. People born abroad to American parents are counted as natives. The census and the ACS survey all residents regardless of legal status, so the immigrant counts include unauthorized immigrants. I measure the size of each local market by its working-age population: everyone aged 16 to 64, whether or not they work. The immigrant counts use the same age range. People born in Puerto

[1] The key IPUMS variables are birthplace (`bpl`), citizenship (`citizen`), age (`age`), and educational attainment (`educ`).

Rico and other U.S. territories are not counted as immigrants, and Alaska and Hawaii are left out of the sample.

I divide immigrants into two education groups. *College-educated* immigrants have any college education. *Non-college* immigrants have no college education. I use these two broad groups, rather than finer education categories, to keep the comparison simple, and I use these two labels throughout the paper.

I aggregate individual records to commuting zones using the crosswalks developed by David Dorn and used in Autor and Dorn (2013) and Autor et al. (2019). Commuting zones are clusters of counties with strong commuting ties (Tolbert and Sizer, 1996), and the 722 zones cover the lower 48 states (Figure 1). Because a commuting zone approximates a local labor market, it also approximates the area in which workers search for housing, so spillovers between neighboring counties stay within the unit of analysis. Including commuting zones that receive few immigrants is also useful, because they serve as a comparison for zones that receive many.



Figure 1: U.S. Commuting Zones

3.2 Housing Data

3.2.1 Zillow Data

I use two series from Zillow Research. The Zillow Observed Rent Index (ZORI) measures typical asking rents. Zillow constructs it as a repeat-rent index, which compares rents for the same units over time, and weights it to reflect the composition of the rental housing stock. The dollar level of the index is based on listed rents between the 35th and 65th percentiles of the local distribution, which excludes the cheapest and most expensive listings (Zillow Research, 2024b). The Zillow Home Value Index (ZHVI) measures the typical home value in a region. It is based on Zillow's estimated market values (Zestimates) for all homes, not only for the homes that sold in a given year (Zillow Research, 2024a).

Zillow coverage is incomplete. ZORI begins in 2015 for many areas and expands gradually, so the rent sample is small. Figure A1 shows that ZORI covers commuting zones in California, Florida, and along the East Coast well, but misses many zones in the interior, including some with large immigrant populations, such as parts of Texas and the northern border. The covered markets are larger and denser than the typical commuting zone, so they may account for a disproportionate share of the identifying variation. ZHVI covers more commuting zones and years than ZORI.

3.2.2 Fair Market Rents and the FHFA House Price Index

To cover all commuting zones and years, I also use Fair Market Rents (FMR) from the U.S. Department of Housing and Urban Development (HUD) and the House Price Index (HPI) from the Federal Housing Finance Agency (FHFA).

HUD sets Fair Market Rents to determine payment standards for the Housing Choice Voucher program and other federal rental assistance. An FMR measures *gross rent* (contract rent plus tenant-paid utilities) for a standard-quality rental unit occupied by a recent mover, excluding public housing and substandard units. In most areas, the FMR is set at the 40th percentile of this distribution; a small number of metropolitan areas use the 50th percentile (U.S. Department of Housing and Urban Development, 2024). FMRs therefore describe the lower-middle part of the rental market, where many immigrants rent. Saiz (2007) also uses FMRs as his main rent measure, from a period when they were set at the 45th percentile. FMRs are not designed for immigrant households, but they are still useful here: if immigration raises the demand for rental housing, rents rise throughout the market, including the units used to compute FMRs.

I measure house prices with the FHFA House Price Index. The FHFA uses a repeat-sales method: it compares the prices of the same homes across transactions, so changes in the index reflect price growth for homes of constant quality. I use the annual county-level all-transactions index, which combines purchase prices with appraisal values from refinance mortgages (Bogin et al., 2019; Federal Housing Finance Agency, 2024). The index is built from mortgages purchased or securitized by Fannie Mae and Freddie Mac. It is normalized to 100 in 2000, so it measures cumulative price growth relative to 2000 rather than price levels. I aggregate county indexes to commuting zones using population weights.

The FHFA index and ZHVI can move differently even over the same period. Suppose the typical home was worth \$100,000 in 2000 and \$120,000 in 2005. If the price of a constant-quality home rises to \$125,000 in 2006, the FHFA HPI rises from 120 to 125, a log change of about 4.1 percent. If, over the same year, new construction and upgrades raise the typical Zillow home value to \$130,000, ZHVI rises by about 8.0 percent. The FHFA index isolates price growth for a fixed set of homes, while ZHVI also reflects changes in the composition and quality of the housing stock. The two measures therefore provide complementary views of the housing market.

I express all dollar values (FMR, ZORI, and ZHVI) in 2005 dollars by dividing by the Consumer Price Index relative to its 2005 level (2005 = 195.3). Deflation does not affect the regression estimates. The outcomes are annual log changes, and a national price deflator subtracts the same amount from every commuting zone in a given year, which the division-by-year fixed effects absorb.

4 Empirical Strategy

4.1 Estimating Equation

For commuting zone c in year t, I estimate

$$\Delta y_{ct} = \beta \, \Delta I_{ct} + \alpha X_{ct_0} + \gamma_{dt} + \varepsilon_{ct}, \tag{1}$$

where Δy_{ct} is the annual log change in FMR, HPI, ZORI, or ZHVI, and

$$\Delta I_{ct} = \frac{M_{ct} - M_{c,t-1}}{L_{c,t-1}} \tag{2}$$

is the change in the immigrant population between years $t-1$ and t, divided by the working-age population in the previous year. Here M_{ct} is the number of immigrants living in commuting zone c in year t, and $L_{c,t-1}$ is the zone's population aged 16 to 64 in year $t-1$. The numerator is a net change: it equals immigrants who arrive in the commuting zone minus immigrants who leave. I refer to ΔI_{ct} as the *net immigrant inflow*. This is the measure used in the census-based specifications of Saiz (2007), except that I divide by the working-age population rather than the total population.

Because the outcome is a log change and ΔI_{ct} is a fraction of the working-age population, β is the percentage change in the housing outcome caused by a net immigrant inflow equal to one percent of the working-age population.[2]

Since the model is estimated in annual changes, β measures a short-run equilibrium response. It combines the increase in housing demand with any construction, vacancy, crowding, tenure, and native mobility responses that occur within the same year.

The term X_{ct_0} contains all the baseline controls: the college share of the working-age population, the manufacturing and construction employment shares, and the sum of origin shares, each measured in the base year t_0 of the instrument and interacted with year indicators; α denotes their coefficients. Division-by-year fixed effects, γ_{dt}, absorb shocks common to all commuting zones in a census division in a given year, such as the regional timing of the housing boom and bust. Section 4.4 describes the controls and standard errors.

[2] For example, if $\beta = 1$, an inflow equal to one percent of the working-age population raises a ZHVI of \$300,000 by about \$3,000, a ZORI of \$1,500 per month by about \$15, and an HPI of 200 by about 2 index points.

4.2 Why OLS Is Biased

OLS estimates of equation (1) are likely biased for three reasons. First, immigrants choose where to live. They move to commuting zones with strong job growth and established immigrant networks, and housing demand in those zones would grow even without new immigrants. Large cities such as New York, Los Angeles, and Chicago attracted many immigrants because of their economic opportunities, and the same opportunities raise housing demand. This sorting biases OLS upward. Second, housing costs also influence where immigrants settle. If fast rent growth discourages new immigrants from moving to a commuting zone, OLS is biased downward. Third, annual immigrant counts come from survey samples and contain sampling error, which biases OLS toward zero. Because these biases work in different directions, the sign of the overall bias is not known in advance. An instrument that shifts immigrant inflows for reasons unrelated to local housing demand addresses all three problems. The instrument I use builds on immigrant network effects, described next.

4.3 The Shift–Share Instrument

Immigrants tend to settle where earlier immigrants from the same country live, because established communities offer job contacts, housing information, and a shared language (Altonji and Card, 1991; Card, 2001). These network effects are part of the endogeneity problem, because networks grew in places whose economies attracted immigrants. The shift–share instrument uses the network effects while avoiding this problem: it predicts where new immigrants go from where earlier immigrants settled years or decades before the study period, so the prediction does not depend on current local conditions. It asks how much each commuting zone's immigrant population would grow if each origin group in the zone grew at its national rate.

Let b index country of birth. The instrument is

$$\Delta z_{ct} = \sum_{b} s_{bct_0}\, g_{bt}, \qquad s_{bct_0} = \frac{M_{bct_0}}{L_{ct_0}}, \qquad g_{bt} = \frac{M_{bt} - M_{b,t-1}}{M_{b,t-1}}, \tag{3}$$

where the share s_{bct_0} is the number of immigrants from origin b living in commuting zone c in base year t_0, divided by the zone's working-age population in that year, and the shock g_{bt} is the national growth rate of the immigrant population from origin b between $t-1$ and t. I construct the instrument separately with $t_0 = 1980$, 1990, and 2000.

For the education groups, I keep the historical shares fixed and replace the national shocks with the growth of the college-educated or non-college immigrant population from each origin:

$$\Delta z^{C}_{ct} = \sum_{b} s_{bct_0}\, g^{C}_{bt}, \qquad \Delta z^{N}_{ct} = \sum_{b} s_{bct_0}\, g^{N}_{bt}. \tag{4}$$

The shares include immigrants of all education levels, because new immigrants of every education level are drawn to places where earlier immigrants from their country settled.

The first stage and the reduced form are

$$\Delta I_{ct} = \nu \, \Delta z_{ct} + \kappa X_{ct_0} + \gamma_{dt} + u_{ct}, \tag{5}$$

$$\Delta y_{ct} = \pi \, \Delta z_{ct} + \mu X_{ct_0} + \gamma_{dt} + v_{ct}. \tag{6}$$

The reduced form links predicted immigration directly to housing outcomes without dividing by an estimated first-stage coefficient, so it remains informative when a first stage is weak. Table A4 reports reduced-form estimates with 1980 shares.

Identifying assumption. Equation (3) is a Bartik-style instrument: a sum over origins of base-year exposure shares multiplied by national shocks (Bartik, 1991). Goldsmith-Pinkham et al. (2020) show that the 2SLS estimator with such an instrument is equivalent to a GMM estimator that uses the historical shares as instruments, with the national shocks determining how the shares are combined. Under this *share-based* view, the identifying assumption is that, after conditioning on the controls and fixed effects, a commuting zone's historical origin shares affect its housing growth in 2005–2019 only through the immigrants they attract, and not through any other, unobserved channel. Borusyak et al. (2022) develop an alternative *shock-based* view, in which identification comes from many national shocks that are as good as randomly assigned, while the shares may be endogenous. I adopt the share-based view. The national inflows in my setting are dominated by a few origins, especially Mexico (Section 5.3), so it is not plausible to treat them as many independent random shocks. The share-based view instead lets me compare commuting zones with different historical origin compositions, and therefore different exposure to common national changes, without assuming that the national inflows themselves are random. The credibility of the estimates therefore depends on the historical shares. This is why I test for pre-trends, show how the shares relate to baseline characteristics, and report Rotemberg weights (Section 5). The historical timing of the shares alone does not establish their validity.

Why I control for the sum of shares. Because each share divides an origin's immigrants by the zone's total working-age population, the shares do not sum to one. They sum to the zone's total immigrant share, $S_{ct_0} = \sum_b s_{bct_0}$. The instrument therefore mixes two sources of variation: *which* origins a zone is exposed to, and *how many* immigrants it had in total. A zone with many immigrants receives a large predicted inflow even if every origin grows at the same national rate, because $\sum_b s_{bct_0} \, \bar{g}_t = S_{ct_0} \, \bar{g}_t$. Borusyak et al. (2022) show that with such incomplete shares the instrument picks up variation in the sum of shares, and they recommend controlling for the sum of shares interacted with period effects. The same control is useful under the share-based view. Without it, part of the identifying variation comes from the overall size of a zone's immigrant population, which is the characteristic most likely to reflect persistent local growth. With it, identification comes from differences in the origin mix. I include the base-year sum of shares, interacted with year indicators, in every controlled specification for this reason.

4.4 Controls and Inference

Baseline controls. A central concern is confounding: persistent local conditions may have influenced both historical immigrant settlement and later housing growth. Commuting zones with strong job markets may have attracted many immigrants in the past and also experienced faster housing demand growth for reasons unrelated to current immigration. If the historical shares reflect such conditions, the instrument is correlated with unobserved determinants of housing growth. To address the observable part of this concern, X_{ct_0} includes the base-year college share of the working-age population, the manufacturing and construction employment shares, and the sum of origin shares, each interacted with year indicators. The sum of shares is rebuilt for each base year. The interactions allow each characteristic to have a different association with housing growth in each year. Using base-year values avoids controlling for characteristics that immigration itself may change. This approach follows Sharpe (2019), who shows that controlling for initial city characteristics reduces the bias of shift–share estimates of immigration's effect on rents.

Land-use regulation. The main specification does not control for land-use regulation. The Wharton Residential Land Use Regulatory Index (WRLURI) that I use is the 2018 version (Gyourko et al., 2021), which was measured after most of the study period. If regulation responded to earlier price growth, controlling for it would remove part of the effect I want to measure. I therefore use WRLURI only to study how the effect of immigration varies with regulation, in an interaction model (Section 6.8).

Inference. I cluster standard errors by commuting zone, which allows for arbitrary correlation of the errors over time within a zone. Adão et al. (2019) show that conventional standard errors can be too small in shift–share designs, because regions with similar shares have correlated residuals. Their correction, like the shock-level inference of Borusyak et al. (2022), is derived for identification from the shocks. Because my identification comes from the shares, I report standard errors clustered by commuting zone.

Education groups. I estimate the model separately for all immigrants, college-educated immigrants, and non-college immigrants. Although the two education groups add up to all immigrants, their separately estimated coefficients do not add up to the total effect. The two education instruments also share the same historical shares and differ only in their national shocks, with Mexico dominating both, so they are likely correlated. When each group is estimated alone, its coefficient may partly reflect the other group's inflows. I therefore also estimate the model with both inflows as endogenous variables, instrumented jointly with the two education-specific instruments, and test whether the two coefficients are equal (Section 6.4).

5 Assessing the Instrument

This section evaluates the shift–share instrument in four steps. Section 5.1 asks whether commuting zones with different historical exposure were already on different housing trends before the study period. Section 5.2 describes the historical shares and how they relate to baseline characteristics (Tables A1 and A2), which motivates the controls. Section 5.3 describes the national shocks. Section 5.4 evaluates the strength of the first stage. Section 6.10 later identifies the origins that drive the estimates. No single check validates the design; together, they show how credible the estimates are and which places they describe.

5.1 Pre-Trends

Figure 2 plots, for each base year, the coefficient from regressing FMR growth over three periods on the shift–share instrument, together with 95 percent confidence intervals. The pattern differs sharply across base years.

With 2000 shares, the instrument does not predict rent growth in the earlier periods. The coefficients for 1980–1990 (about 0.2) and 1980–2000 (about 0.1) are close to zero, and their confidence intervals comfortably include zero. In the study period, 2005–2019, the coefficient rises to about 1.3 and is statistically significant.

With 1990 shares, the earlier coefficients are negative but small (about -0.2 and -0.4) and statistically insignificant. The study-period coefficient is about 1.5, with a confidence interval that just includes zero.

With 1980 shares, the earlier coefficients are about -1.0 and -0.9, and their 95 percent confidence intervals lie entirely or almost entirely below zero. Commuting zones with more exposure under 1980 shares therefore had slower rent growth in the 1980s and 1990s, and the relationship turns positive (about 1.0, not significant) only in the study period. This pattern suggests that the 1980 shares are correlated with long-run rent trends that are unrelated to immigration in the study period. The switch from a negative to a positive relationship could also reflect mean reversion rather than an effect of immigration.

None of the three base years allows a strict pre-trend test, because the census data begin in 1980 and both earlier windows start in 1980. With 1980 shares, both windows begin in the base year itself. With 1990 shares, the 1980–2000 window extends past the base year. With 2000 shares, both windows end at or before the base year, so the 2000 shares partly reflect where immigrants settled during those windows. The test therefore asks whether exposure predicts rent growth in earlier periods, not strictly before the base year. On this test, the 1980 shares raise the most concern, while the 1990 and 2000 shares show no significant relationship with earlier rent growth. Because the aim of this paper is to compare how the shift–share instrument behaves across base years, I report results for all three.

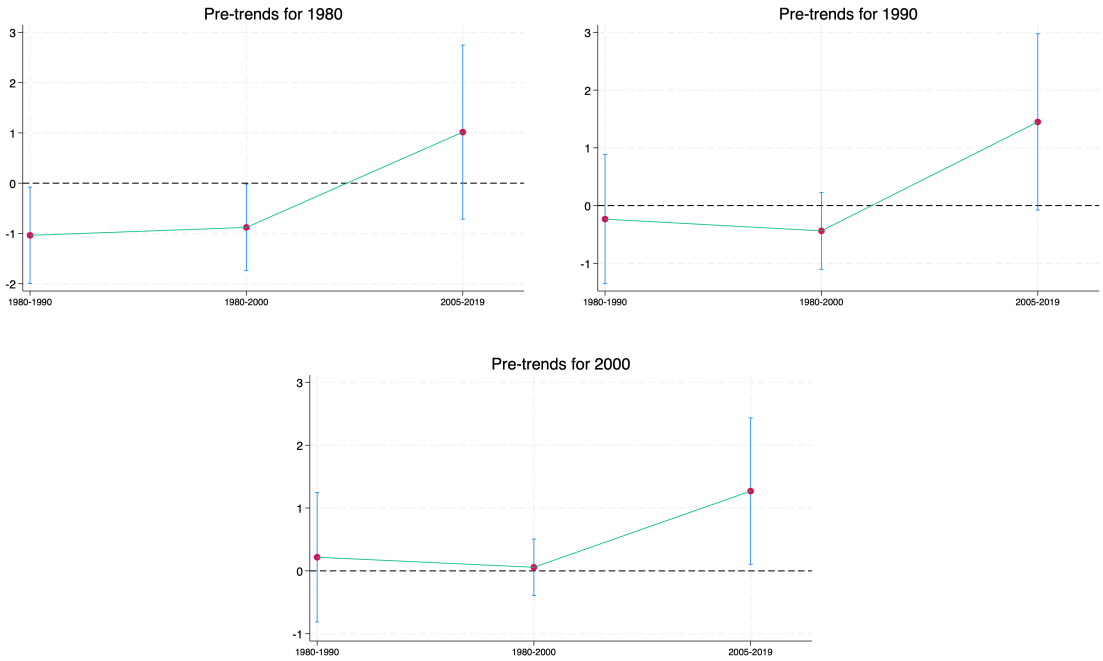


Figure 2: Pre-Trend Tests Using 1980, 1990, and 2000 Shares

Notes: Each panel plots ρ_p from equation (7) for the indicated base year, with 95 percent confidence intervals based on standard errors clustered by commuting zone. The outcome is the log change in FMR over each period.

The test. Each point in Figure 2 is the coefficient ρ_p from

$$\Delta y_{cp} = \rho_p \, \Delta z_{cp} + \alpha_p X_{ct_0} + \gamma_{dp} + \eta_{cp}, \tag{7}$$

estimated separately for each period $p \in \{1980\text{–}1990,\ 1980\text{–}2000,\ 2005\text{–}2019\}$. Here Δy_{cp} is the log change in FMR in commuting zone c over period p, Δz_{cp} is the shift–share instrument, and ρ_p is the coefficient on the instrument, not the change in rents itself. In each period, the instrument combines the base-year shares with the national growth of each origin group over that same period. For 1980–1990, for example, $\Delta z_{cp} = \sum_b s_{bct_0} \, g_{bp}$, where g_{bp} is the growth of origin b's national immigrant population between 1980 and 1990. Because the census data begin in 1980, no period precedes 1980. The controls are the same base-year characteristics as in the main specification: the college share, the manufacturing and construction shares, and the sum of origin shares. I use FMR because it is the only outcome available for all commuting zones in every period.

Why this test matters. The shift–share estimator compares commuting zones with different historical exposure to immigration. This comparison is valid only if high- and low-exposure zones would have followed similar housing trends without the immigration of the study period. The pre-trend test checks the part of this assumption that can be observed. If the instrument predicts rent growth in a period when it should not matter, the shares are probably correlated with persistent local conditions, such as industry or amenity trends, that also drive rents, and 2SLS would attribute those trends to immigration. The

test is the analogue of the pre-trend check in a difference-in-differences design with a continuous treatment: as Goldsmith-Pinkham et al. (2020) show, each origin share acts as a separate measure of treatment intensity. Passing the test does not prove that the shares are exogenous, because trends could change exactly when the study period begins, but failing it is strong evidence against the design. Because each period uses its own national shocks, the instrument also predicts the immigrants who arrived during that period, so an earlier-period coefficient combines any pre-existing trend with any effect of immigration in that period. A clearly negative coefficient, as with 1980 shares, more likely reflects a trend, since Section 2 predicts that immigration raises rents. A coefficient close to zero, as with 2000 shares, is consistent with no trend, but it could also reflect a positive effect of immigration offset by a negative trend. The test therefore indicates how much weight to place on each base year, but it does not establish that any set of shares is valid.

5.2 Historical Settlement Shares

The choice of base year involves a trade-off. Older shares are less likely to reflect conditions that also drive housing markets in 2005–2019, but they predict current settlement less well. More recent shares capture current immigrant networks better and give a stronger first stage, but they are more likely to reflect persistent local conditions. Comparing three base years shows whether the findings depend on one particular historical distribution of immigrants. It also helps interpret differences across specifications, although comparisons across base years alone cannot identify the economic mechanisms behind the housing response.

1980 Shares (U.S. immigrant share, ages 16–64 = 6.2%)					1990 Shares (U.S. immigrant share, ages 16–64 = 7.8%)					2000 Shares (U.S. immigrant share, ages 16–64 = 11.1%)				
Origin	Mean	Min	Max	SD	Origin	Mean	Min	Max	SD	Origin	Mean	Min	Max	SD
Mexico	0.72	0.00	21.75	2.31	Mexico	1.27	0.00	28.54	3.22	Mexico	2.36	0.00	33.73	4.19
Canada	0.25	0.00	6.19	0.41	Germany	0.19	0.00	1.44	0.17	Canada	0.22	0.00	3.74	0.30
Germany	0.22	0.00	1.88	0.18	Canada	0.19	0.00	4.72	0.34	Germany	0.17	0.00	1.42	0.15
United Kingdom	0.13	0.00	0.74	0.12	United Kingdom	0.12	0.00	0.65	0.12	Philippines	0.15	0.00	4.56	0.32
Philippines	0.06	0.00	2.55	0.16	Philippines	0.10	0.00	3.99	0.27	China	0.13	0.00	4.86	0.30
Italy	0.06	0.00	2.00	0.15	China	0.09	0.00	3.76	0.21	United Kingdom	0.13	0.00	0.88	0.13
Cuba	0.06	0.00	16.55	0.63	Vietnam	0.07	0.00	2.47	0.16	India	0.13	0.00	2.57	0.21
China	0.05	0.00	2.49	0.13	Korea	0.07	0.00	1.19	0.11	Vietnam	0.12	0.00	4.43	0.26
Korea	0.05	0.00	0.86	0.08	India	0.07	0.00	1.22	0.11	Korea	0.09	0.00	1.46	0.13
Vietnam	0.04	0.00	0.56	0.07	Cuba	0.05	0.00	16.58	0.63	El Salvador	0.08	0.00	2.69	0.19

Notes: For the ten largest origins in each base year, the table reports the mean, minimum, maximum, and standard deviation across the 722 commuting zones of s_{bct_0}, the number of immigrants from origin b divided by the commuting zone's working-age population, in percent. The column headers report the immigrant share of the U.S. working-age population (ages 16–64) in each base year.

Table 1: Immigrant Shares of the Commuting-Zone Working-Age Population (%), Ten Largest Origins

Table 1 summarizes the shares for the ten largest origins in each base year. The immigrant share of the U.S. working-age population rose by 4.9 percentage points between 1980 (6.2 percent) and 2000 (11.1 percent), and it grew faster in the 1990s (from 7.8 to 11.1 percent) than in the 1980s. Immigration policy contributed to the faster growth in the 1990s. The Immigration Act of 1990 raised the annual worldwide limit on immigrant visas from 270,000 to a flexible limit of 700,000 for fiscal years 1992–1994 and 675,000 thereafter,

expanded employment-based visas, and created the Diversity Visa program.[3] Its main provisions took effect in fiscal year 1992, after the 1990 census. The act can therefore help explain the change in shares between 1990 and 2000, but not the 1990 shares themselves. It also diversified the origins of new immigrants and increased the admission of skilled workers.

Mexico has the largest average share in all three base years. In 1980, Mexican immigrants made up 0.72 percent of the working-age population in the average commuting zone, but more than 20 percent in some commuting zones along the Texas border. The average Mexican share rose to 1.27 percent in 1990 and 2.36 percent in 2000, and the maximum reached 33.7 percent in 2000. Canada ranks second in 1980 and 2000, with average shares of 0.25 and 0.22 percent, and third in 1990 (0.19 percent). The maximum Canadian share fell from 6.19 percent in 1980 to 3.74 percent in 2000. Canadian immigrants are concentrated in Washington and in northern border states such as Maine, New York, Vermont, and New Hampshire. German-born immigrants follow a similar pattern, with the maximum share falling from 1.88 percent in 1980 to 1.42 percent in 2000, and their largest concentrations are in Texas, Oklahoma, and Colorado.[4]

Cuba has a small average share but the most extreme concentration. In 1980, Cuban immigrants made up 16.6 percent of the working-age population in one commuting zone, the Miami commuting zone. This concentration reflects the Cuban exile waves of the 1960s and 1970s, not the Mariel boatlift. The 1980 census was taken on April 1, 1980, while the boatlift, which brought about 125,000 Cubans to the United States, took place between April and October 1980. The Mariel arrivals therefore appear in the 1990 shares, in which the maximum Cuban share is again about 16.6 percent.

Immigration from Asia also grew in importance. The average share of immigrants from China rose from 0.05 percent in 1980 to 0.13 percent in 2000, from India from 0.07 percent in 1990 to 0.13 percent in 2000, and from Vietnam from 0.04 percent in 1980 to 0.12 percent in 2000. These origins moved up the rankings over the two decades, and their immigrants settled mainly in California, New York, and Illinois.

The shares include immigrants of all education levels, but the education of the immigrant population changed over the base years (Figure 3). The share of immigrants without a high school diploma fell from 52.4 percent in 1980 to 39.0 percent in 2000. The share with exactly four years of college, equivalent to a bachelor's degree, nearly doubled over these two decades, from 5.8 to 11.5 percent, and the share with at least four years of college rose from 12.7 to 19.7 percent. The share in the three college categories of Figure 3 rose from 25.2 percent in 1980 to 34.4 percent in 1990 and 33.2 percent in 2000. A more educated

[3] Immigration Act of 1990, Pub. L. No. 101-649, 104 Stat. 4978.

[4] These concentrations likely reflect large U.S. military installations, such as Fort Hood (now Fort Cavazos) in Texas, Fort Sill in Oklahoma, and Fort Carson in Colorado. German-born spouses and children of service members who were stationed in Germany would be counted as German-born residents near these bases.

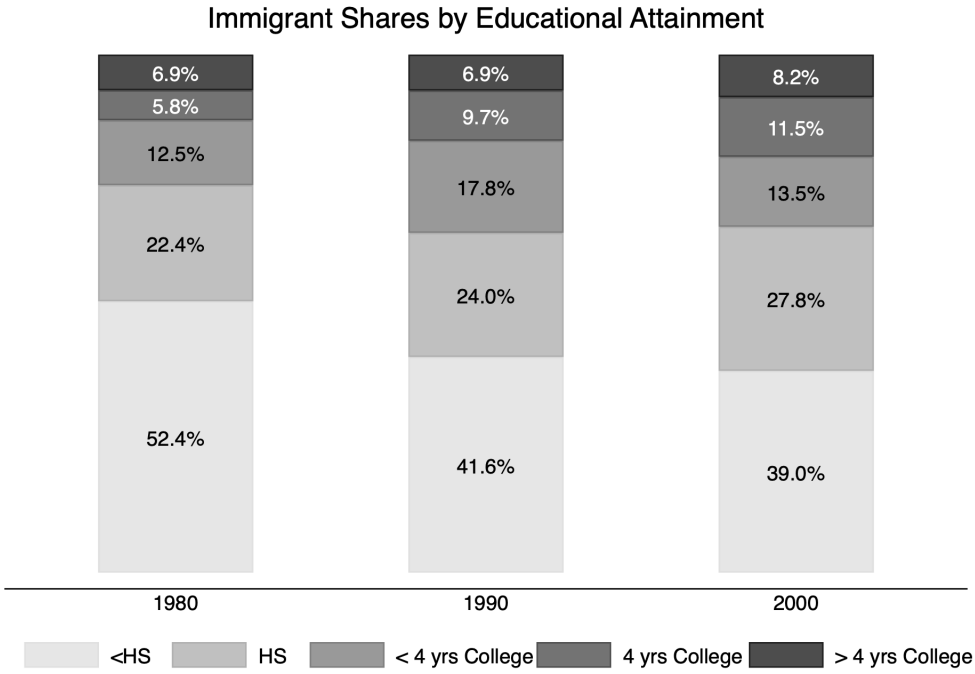


Figure 3: Educational Attainment of Immigrants, 1980, 1990, and 2000

immigrant population has greater earning potential, which may allow immigrants to enter the housing market as buyers sooner.

Immigrants are also much more geographically concentrated than natives. Figure A2 shows that immigrants make up a small share of the working-age population in most commuting zones and a large share in only a few, while native shares are spread much more evenly. College-educated immigrants are a small share of the working-age population in most commuting zones, with a long tail of a few zones where their share is high. Tables A1 and A2 show that the origin shares are correlated with baseline college, manufacturing, and construction shares, and with the total immigrant share. For example, in 1980 the Mexican share is lower in commuting zones with larger manufacturing shares. These correlations are why the main specification controls for these characteristics.

5.3 National Shocks

The national shocks are the annual growth rates of each origin's immigrant population. Figure 4 plots the national immigrant population aged 16 to 64 from the largest origins in each year from 2005 to 2019, for all immigrants and for each education group. These are national counts, shown to give a sense of the size of each origin group.

Mexico is by far the largest origin. The Mexican-born population is between 7 and 8 million throughout the period and changes little from year to year, rising slightly until 2014 and declining after 2016. Mexico also dominates the non-college group, at about 6 million, with a decline after 2016. India, China, the Philippines, and El Salvador are

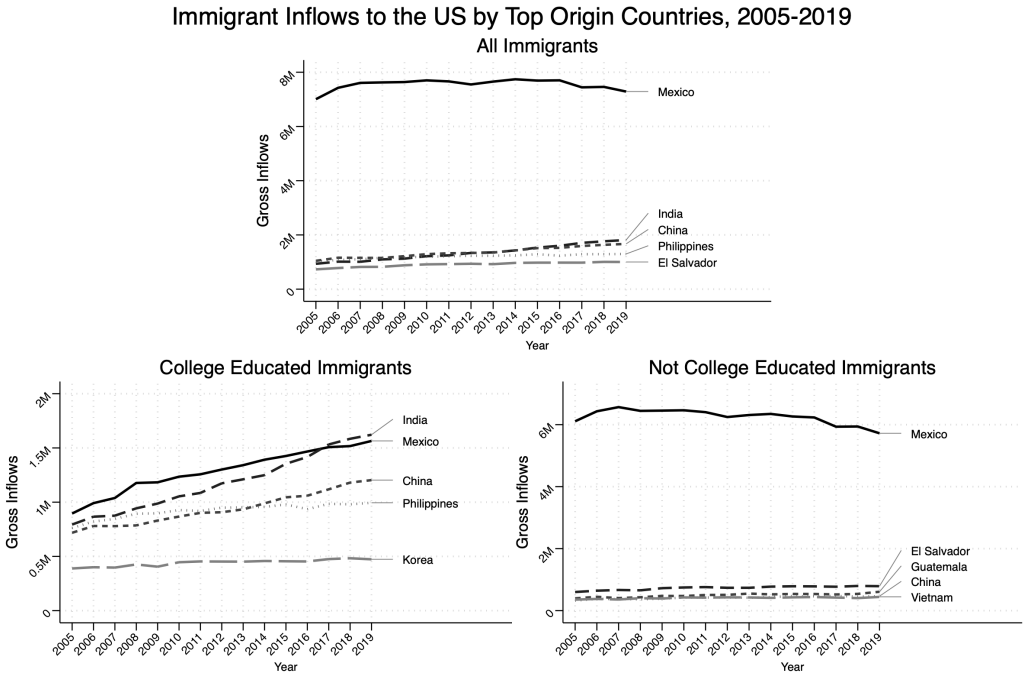


Figure 4: National Immigrant Population (Ages 16–64) by Origin and Education, 2005–2019

the next largest origins among all immigrants; they are far smaller than Mexico but grow steadily. Among non-college immigrants, El Salvador, Guatemala, China, and Vietnam follow Mexico.

The college group looks different. Mexico and India are the largest origins, each growing from under 1 million in 2005 to about 1.6 million in 2019, and India overtakes Mexico in 2017. China and the Philippines follow at about 1 million each, and the Korean population is stable at about 0.5 million. The college group is much smaller than the non-college group for every origin.

Two features of the shocks matter for identification. First, the largest current origins (Mexico, China, India, the Philippines, and El Salvador) match the 2000 shares more closely than the 1980 or 1990 shares, which helps explain why the 2000-share instrument has the strongest first stage. Second, because the Mexican population is large and grows slowly and steadily, much of the instrument comes from the same places receiving similar shocks year after year. When the origin mix of inflows is this stable, the instrument can pick up both the short-run effect of current inflows and the continuing adjustment to past inflows into the same places (Jaeger et al., 2018).

5.4 First Stage

Table 2 reports estimates of the first-stage equation (5) for each base year, without controls (FE) and with baseline controls (FE + C). The sample has 10,108 observations: 722 commuting zones over 14 annual changes. The coefficient ν measures how much the

actual net inflow rises when the predicted inflow rises by one percentage point.

The instrument is strong for all immigrants in every base year. With 1980 shares and no controls, $\nu = 1.45$: a one-percentage-point increase in the predicted inflow is associated with a 1.45-percentage-point increase in the actual net inflow. With controls, the coefficient rises to 1.79. The first stage is strongest with 2000 shares ($F = 84$ with controls) and 1990 shares ($F = 80$), compared with $F = 33$ for 1980 shares. Without controls, the coefficient falls as the base year approaches the study period, from 1.45 (1980) to 1.11 (1990) and 0.76 (2000). With controls, it is similar for 1980 and 1990 (1.79 and 1.80) and lower for 2000 (1.38).

A first-stage coefficient above one is not surprising given how the instrument is built. The instrument applies national growth rates to base-year shares and the base-year working-age population, while the regressor is the actual change divided by last year's working-age population. If every origin group in a commuting zone grew at its national rate, the actual inflow would equal the predicted inflow multiplied, origin by origin, by the national growth of that origin group since the base year ($M_{b,t-1}/M_{bt_0}$) and by the change in the zone's working-age population ($L_{ct_0}/L_{c,t-1}$). The main origin groups, Mexico in particular, grew several-fold between 1980 and the late 2000s, so the instrument understates actual inflows by a factor that is larger the older the base year. This is consistent with the decline in the uncontrolled coefficient from 1.45 with 1980 shares to 0.76 with 2000 shares. Network effects may add to this: new immigrants may concentrate in established communities more than in proportion to their size.

For the education groups, the national shocks are computed separately for college-educated and non-college immigrants, while the historical shares are the same. With 1980 shares and controls, a one-percentage-point increase in the predicted inflow is associated with increases of 0.58 and 0.56 percentage points in the actual college and non-college inflows. The non-college first stage is adequate in every base year, with F-statistics of 14, 16, and 31 with controls. The college first stage is weak with 1980 and 1990 shares once controls are added ($F = 4$ and 3), but strong with 2000 shares ($F = 42$, coefficient 0.60). With F-statistics as low as 3 or 4, 2SLS estimates are centered near the OLS estimates and conventional t-tests reject too often, so I treat the college estimates based on 1980 and 1990 shares as uninformative.

The 2000 shares likely perform best because they are closest in time to the study period and better reflect the networks that direct current immigrants. The origins whose shares grew most by 2000, including Mexico, the Philippines, China, El Salvador, and Vietnam, are also among the largest origins of current immigrants (Figure 4). Because the 2000 shares give the strongest first stage, they are the preferred specification on this criterion. The aim of this paper, however, is to compare how the shift–share instrument behaves across base years, so I report and compare the results for all three.

	1980 Shares		1990 Shares		2000 Shares	
	FE	FE + C	FE	FE + C	FE	FE + C
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: All Immigrants						
Predicted inflow (Δz_{ct})	1.4516	1.7937	1.1073	1.8039	0.7584	1.3841
	(0.2498)	(0.3100)	(0.2036)	(0.2016)	(0.1055)	(0.1509)
F-statistic	34	33	30	80	52	84
Panel B: College-Educated Immigrants						
Predicted inflow (Δz^{C}_{ct})	0.8730	0.5752	0.5912	0.4914	0.4001	0.6036
	(0.1630)	(0.2985)	(0.1425)	(0.3063)	(0.0595)	(0.0927)
F-statistic	29	4	17	3	45	42
Panel C: Non-College Immigrants						
Predicted inflow (Δz^{N}_{ct})	0.5726	0.5609	0.4222	0.5112	0.3002	0.4298
	(0.1940)	(0.1516)	(0.1166)	(0.1278)	(0.0577)	(0.0768)
F-statistic	9	14	13	16	27	31
Observations	10,108	10,108	10,108	10,108	10,108	10,108

Notes: The dependent variable is the net immigrant inflow ΔI_{ct} for the indicated group. All regressions include division×year fixed effects and use analytic weights based on the 1980 working-age population. C denotes controls: the base-year college share, manufacturing share, construction share, and sum of origin shares, each interacted with year indicators. Standard errors clustered by commuting zone are reported in parentheses. The F-statistic is the squared ratio of the first-stage coefficient to its clustered standard error; with one instrument and one endogenous variable, it equals the Kleibergen–Paap Wald F-statistic (Kleibergen and Paap, 2006) and the effective F-statistic of Montiel Olea and Pflueger (2013).

Table 2: First-Stage Estimates by Base Year

6 Results

6.1 Summary Statistics

Table 3 reports summary statistics for the four housing outcomes over 2005–2019. FMR and HPI cover nearly all commuting zones, while the Zillow series cover fewer zones and expand over time. The average Zillow rent is $848 per month, compared with $729 for FMR. This gap mostly reflects the different samples: ZORI covers mainly large, dense markets and begins in 2015, while FMRs cover nearly all commuting zones from 2005 onward. The lower FMR target, the 40th percentile of rents, adds to the gap, while the inclusion of utilities in FMRs works in the opposite direction. The minimum Zillow rent ($424) is also higher than the minimum FMR ($388), but the maximum FMR is higher, perhaps because FMRs cover more commuting zones and years. The average HPI is 144, so prices for a constant set of homes averaged 44 percent above their 2000 level. This comparison is in nominal terms, because I do not deflate the HPI. The average Zillow home value is $125,979, ranging from about $25,000 to $1.3 million.

Figures A3 and A4 show how the housing measures moved over the study period. Although the FHFA HPI and ZHVI are built differently, both fall sharply around the Great Recession and then recover gradually (Figure A3). In annual log changes (Figure A4), home prices

	Mean	SD	Min	Max
FMR ($ per month)	729.09	203.00	387.92	2,575.34
FHFA HPI (2000 = 100)	143.62	27.57	77.05	403.99
ZORI ($ per month)	848.00	248.68	423.96	2,372.87
ZHVI ($)	125,979.30	79,029.34	25,415.32	1,338,685.00

Notes: Unweighted statistics across commuting-zone–year observations, 2005–2019. Dollar values are in constant 2005 dollars. HPI is an index equal to 100 in 2000.

Table 3: Summary Statistics for Housing Outcomes

respond sharply around 2009, while rents decline more gradually and reach their trough after 2010. This lag may reflect the fact that most residential leases are signed for a year.

Table A3, together with Figures A5 and A6, compares the four measures in log levels and log changes. If the measures captured the same housing market conditions, they would be highly correlated in both. They are highly correlated in levels: places with high rents also have high home values. In annual changes, however, the correlations are much weaker. The two home-value measures are the most closely related, with weighted correlations of 0.43 in log changes and 0.74 in log levels. Home values tend to move together because both measure the price of housing as an asset, which responds to interest rates, credit conditions, and expected appreciation. Rents are the current price of housing services. They adjust more slowly because of lease contracts and respond less directly to financing conditions, so rent measures are only weakly correlated with home-value measures in annual changes. The FMR and ZORI series are almost uncorrelated in annual changes (0.07) even though their levels are highly correlated (0.92), which shows that the two rent series do not record the same short-run changes. Differences in geographic coverage add to these gaps, because the Zillow series cover a narrower set of markets.

6.2 Main Results

Table 4 reports 2SLS estimates of equation (1) with baseline controls for each base year. Each coefficient is the percentage change in the housing outcome caused by a net immigrant inflow equal to one percent of the working-age population. Because the four outcomes cover different commuting zones and years, the number of observations differs across columns.

Immigration raises market rents and home values with all three base years. A net inflow equal to one percent of the working-age population raises Zillow home values by 3.3 to 3.7 percent and Zillow rents by 1.4 to 3.6 percent, all statistically significant. The constant-quality FHFA index rises by about 1.5 percent (1.45 to 1.52, also statistically significant). The FMR estimates, between 0.3 and 0.8, are insignificant in every base year. The ZHVI and HPI estimates are very similar across base years. The ZORI estimate is smaller with 1980 shares (1.4) than with 1990 or 2000 shares (2.6 and 3.6). With 2000 shares, which give the strongest first stage, the inflow raises a monthly Zillow rent of $848 by about $31 and a Zillow home value of $125,979 by about $4,100, at the sample means in Table 3.

	FMR (1)	FHFA HPI (2)	ZORI (3)	ZHVI (4)
Panel A: 1980 Shares				
All immigrants	0.2772	1.4868***	1.4050***	3.7271***
	(0.5376)	(0.5383)	(0.3953)	(0.8389)
College-educated immigrants	-0.7201	4.7096**	-2.1450**	6.2552***
	(1.7525)	(2.0073)	(1.0382)	(2.3996)
Non-college immigrants	1.0472	0.4256	4.3851***	4.9391***
	(1.0740)	(1.0033)	(0.8231)	(1.7808)
Panel B: 1990 Shares				
All immigrants	0.7862	1.5217**	2.5995***	3.4287***
	(0.6942)	(0.6117)	(0.5630)	(0.8624)
College-educated immigrants	4.2796	6.3045	-1.3478	3.9284
	(8.1781)	(5.5180)	(1.7714)	(5.4358)
Non-college immigrants	1.9447	1.0916	5.7178***	5.5024***
	(1.3419)	(1.0778)	(1.0213)	(1.8676)
Panel C: 2000 Shares				
All immigrants	0.8051	1.4460**	3.6087***	3.2870***
	(0.6873)	(0.6830)	(0.7764)	(0.7420)
College-educated immigrants	3.4998	1.9334	1.5056	1.4013
	(3.6840)	(1.6325)	(1.2717)	(3.0750)
Non-college immigrants	1.5324	1.6720	6.9581***	5.0423***
	(1.2991)	(1.2349)	(1.5654)	(1.5008)
Observations	10,108	9,126	868	8,095

Notes: Each coefficient comes from a separate 2SLS regression of the annual log change in the housing outcome on the net immigrant inflow ΔI_{ct} of the indicated group, instrumented with the shift–share instrument for that group and base year. All regressions include division×year fixed effects and baseline controls (the base-year college share, manufacturing share, construction share, and sum of origin shares, each interacted with year indicators), exclude WRLURI, and use analytic weights based on the 1980 working-age population. First-stage estimates are in Table 2; the college-educated first stage is weak with 1980 and 1990 shares ($F = 4$ and 3). Standard errors clustered by commuting zone are reported in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table 4: 2SLS Estimates of the Effect of Immigration on Housing Outcomes

Why are the Zillow estimates larger than the FMR and HPI estimates? The four outcomes differ both in what they measure and in where they are observed. ZORI measures current asking rents, while FMRs are based on ACS rents that are two to three years old and then projected forward, so annual changes in FMRs respond slowly to current shocks. The FHFA HPI measures constant-quality prices of homes financed with conforming mortgages, while ZHVI measures the value of the whole housing stock, including new and upgraded homes. The samples also differ: the Zillow series cover fewer commuting zones, mostly large and dense markets where housing supply is less elastic. Section 6.7 holds the sample fixed to separate these explanations.

6.3 Results by Education

Non-college immigrants. Estimated one group at a time, the housing response comes mainly from immigrants without college education. A non-college inflow equal to one percent of the working-age population raises Zillow rents by 4.4 to 7.0 percent and Zillow home values by 4.9 to 5.5 percent across the three base years, all statistically significant. The HPI and FMR estimates are positive in every base year but statistically insignificant. Compared with the uncontrolled estimates in Table A6, the controls lower the ZORI estimates (from 5.1, 7.2, and 9.2 to 4.4, 5.7, and 7.0), leave the ZHVI estimates roughly unchanged, and make the HPI estimates insignificant (from about 2.2 to 2.5 without controls to 0.4 to 1.7 with controls). Non-college immigrants are more closely tied to local industry structure, so the manufacturing and construction controls remove more of the correlation between their predicted inflows and housing growth. Once these components are removed, the effect remains in the Zillow-covered markets but is no longer detectable in the broader HPI and FMR samples. This suggests that the housing response to non-college immigration is concentrated in dense, high-demand markets.

College-educated immigrants. With 2000 shares, where the college first stage is strong ($F = 42$), the college estimates are small relative to their standard errors: 3.5 (standard error 3.7) for FMR, 1.9 (1.6) for HPI, 1.5 (1.3) for ZORI, and 1.4 (3.1) for ZHVI. None is statistically significant. With 1980 shares and controls, the college estimates are large and significant for HPI (4.7) and ZHVI (6.3) and negative for ZORI (-2.1), but the first-stage F-statistic is only 4. With 1990 shares, the first stage is equally weak ($F = 3$), and no estimate is significant. Because 2SLS estimates with weak instruments are centered near the OLS estimates and conventional t-tests reject too often, I do not interpret the college estimates based on 1980 and 1990 shares. The evidence is consistent with no short-run housing effect of college-educated immigration, but the estimates are too imprecise to rule out meaningful effects.

These comparisons rely on separate regressions for each group. Larger and more precise non-college estimates do not by themselves show that the two effects differ. Section 6.4 estimates the two groups jointly and tests this directly.

6.4 Estimating the Two Education Groups Jointly

Because the college-educated and non-college instruments are built from the same historical shares, the separate estimates in Section 6.3 may partly attribute one group's inflows to the other. Appendix Tables A7 and A8 therefore estimate the model with both inflows as endogenous variables, instrumented jointly with the two education-specific instruments. The joint model is

$$\Delta y_{ct} = \beta_C \, \Delta I^C_{ct} + \beta_N \, \Delta I^N_{ct} + \alpha X_{ct_0} + \gamma_{dt} + \varepsilon_{ct}, \tag{8}$$

where ΔI^C_{ct} and ΔI^N_{ct} are the net inflows of college-educated and non-college immigrants, each defined as in equation (2). Both inflows are endogenous, and each has its own first stage that includes both education-specific instruments from equation (4):

$$\Delta I^C_{ct} = \nu_{CC} \, \Delta z^C_{ct} + \nu_{CN} \, \Delta z^N_{ct} + \kappa_C X_{ct_0} + \gamma_{dt} + u^C_{ct}, \tag{9}$$

$$\Delta I^N_{ct} = \nu_{NC} \, \Delta z^C_{ct} + \nu_{NN} \, \Delta z^N_{ct} + \kappa_N X_{ct_0} + \gamma_{dt} + u^N_{ct}. \tag{10}$$

The specification without controls omits the X_{ct_0} terms. Each column of the tables reports β_C and β_N, the p-value from a test of $\beta_C = \beta_N$, and the first-stage F-statistic. With two endogenous variables, the F-statistic is the Kleibergen–Paap rk Wald F-statistic (Kleibergen and Paap, 2006), which measures whether the two instruments are jointly strong enough to identify both coefficients. Unlike the F-statistics in Table 2, it is not the squared t-statistic of a single first-stage coefficient, and it is low when the two instruments move the two inflows in similar proportions.

Without baseline controls. The joint first stage is reasonably strong with 1990 and 2000 shares (F between 16 and 29 for FMR, HPI, and ZHVI) and weaker with 1980 shares (F about 8) and in the ZORI sample (F between 7 and 15). Two patterns stand out (Table A7). First, non-college inflows raise home values. They raise the FHFA index by 2.4 to 3.7 percent and Zillow home values by 4.1 to 6.1 percent in every base year, while college-educated inflows have no significant positive effect on either measure. The difference between the two groups is significant for HPI in all three base years ($p \leq 0.031$) and for ZHVI with 1990 and 2000 shares ($p = 0.003$ and $p < 0.001$). Second, both groups raise Zillow rents. College-educated inflows raise ZORI by 3.1 to 5.8 percent, and non-college inflows by 2.8 to 3.7 percent (significant with 1980 and 2000 shares), and I cannot reject that the two effects are equal in any base year (p between 0.36 and 0.86). Neither group has a significant effect on FMR.

With baseline controls. Adding the baseline controls weakens the joint first stage considerably (Table A8). The joint first-stage F-statistic is about 8 with 1980 shares (2.8 for ZORI), between 1.1 and 1.4 with 1990 shares, and between 2.0 and 4.1 with 2000 shares, well below the rule-of-thumb value of 10. The controlled joint estimates are therefore imprecise and should be read with caution. One result nevertheless holds in

every base year: non-college inflows raise Zillow rents, by 5.3 percent with 1980 shares, 6.2 percent with 1990 shares, and 5.8 percent with 2000 shares, all statistically significant. These estimates are close to the separate non-college estimates in Table 4 (4.4 to 7.0 percent). With 2000 shares, non-college inflows also raise Zillow home values by 5.7 percent, although this estimate is less precise. The equality test does not reject equal effects for ZORI in any base year (p between 0.10 and 0.24), because the college-educated estimates are very imprecise. With 1980 shares, the controlled estimates show the opposite pattern for home values: college-educated inflows raise HPI by 2.6 percent and ZHVI by 5.4 percent, non-college inflows have negative or insignificant effects, and the differences are significant ($p = 0.010$ and 0.042). Given the weak joint first stage and the weak separate college first stage with 1980 shares, I do not place much weight on this result. With 1990 and 2000 shares, no difference between the two groups is significant.

6.5 Magnitudes Compared with the Literature

Saiz (2007) finds that an immigrant inflow equal to one percent of a metropolitan area's population raises average rents and housing values by about one percent. My Zillow estimates are larger: 1.4 to 3.6 percent for rents and 3.3 to 3.7 percent for home values. My FMR point estimates, between 0.3 and 0.8, are closer to Saiz's estimate, which is also based mainly on FMRs, but they are imprecise. Saiz divides immigrant inflows by the total population, while I divide by the working-age population, so the two sets of estimates are not in identical units.

Several differences can explain this gap. First, the outcomes differ. Saiz's main rent measure is the FMR, which was then set at the 45th percentile of rents, while ZORI and ZHVI measure current market rents and home values in large, dense markets. Second, the regressor differs. Saiz's annual measure counts newly admitted legal immigrants, while mine is the net change in the total immigrant population, including unauthorized immigrants and net of departures. Because departures make a net change smaller than the arrivals behind it, the same demand shock yields a larger coefficient per unit of net change; the omission of unauthorized arrivals from Saiz's measure works in the opposite direction, so the net effect of this difference is unclear. Third, the period differs. Saiz studies the 1980s and 1990s, while my period includes the housing bust and a recovery in which housing supply was slow to respond, which may have raised the price response to a given demand shock. Estimates from other countries also exceed Saiz's: Degen and Fischer (2017) find that an inflow equal to one percent of a Swiss region's population raises single-family house prices by about 2.7 percent. My Zillow estimates are at the upper end of this range and should be read as effects for the large, dense markets that Zillow covers.

6.6 The Role of Baseline Controls

Appendix Table A6 reports the same models without baseline controls. Adding the controls raises the all-immigrant estimates for FMR, HPI, and ZHVI in every base year. With

2000 shares, for example, the ZHVI estimate rises from 0.71 to 3.29 and the HPI estimate from −0.18 to 1.45. The ZORI estimates move in the other direction, falling from 2.43 to 1.41 (1980 shares), from 4.02 to 2.60 (1990 shares), and from 4.61 to 3.61 (2000 shares).

These changes indicate omitted-variable bias in the uncontrolled estimates, and its direction depends on each omitted characteristic. Omitting a characteristic biases the estimate downward when the characteristic is negatively correlated with exposure to the instrument but positively correlated with housing growth, and upward when both correlations have the same sign. Table A1 shows that in 1980, the Mexican share, the largest component of the instrument, is lower in commuting zones with higher manufacturing shares and, less precisely, higher college shares; the 2000 Mexican share shows the same signs, although neither is significant (Table A2). If commuting zones with more college graduates experienced faster housing growth over 2005–2019, omitting the college share understates the effect of immigration, and adding it raises the estimate. The manufacturing share works in the opposite direction if commuting zones with more manufacturing experienced slower housing growth: omitting it overstates the effect of immigration, and adding it lowers the estimate. The net effect of the controls therefore depends on which of these biases is larger. Because adding the controls raises the FMR, HPI, and ZHVI estimates, the downward bias appears to dominate in these samples. The increase is largest for ZHVI. Home values capitalize expected future rents, and this forward-looking component may be strongest in the dense, high-skill markets where the controls matter most. The ZORI sample is limited to Zillow-covered markets, where the relationship between baseline characteristics and exposure may differ, so the controls can move the ZORI estimates in the opposite direction.

6.7 Common Zillow Sample

ZORI has the most limited coverage of the four outcomes. Its observations are concentrated in large metropolitan areas, where immigrant populations are larger and housing supply is less elastic (Figure A1). To hold geography fixed, Table A10 re-estimates all four outcomes on the 814 observations for which all four are available, and Table A5 reports first-stage estimates in the Zillow rent sample.

Restricting the sample raises the HPI and ZHVI estimates substantially. For all immigrants with 2000 shares, the HPI estimate rises from 1.4 to 2.9 and the ZHVI estimate from 3.3 to 6.6, while the FMR estimate remains insignificant. The same exposure to immigration therefore moves house prices more in Zillow-covered markets than in the national sample, so the larger Zillow estimates partly reflect *where* Zillow measures housing, not only *what* it measures. The first stage is weaker in this sample, especially for college-educated immigrants (F between 3 and 13 with controls), and conventional significance can be misleading when instruments are weak. These estimates therefore show that the results are sensitive to the sample, but the larger coefficients may partly reflect weaker identification rather than stronger housing responses.

6.8 Land-Use Regulation

To examine whether the effect of immigration varies with the restrictiveness of local land-use regulation, I estimate

$$\Delta y_{ct} = \beta\,\Delta I_{ct} + \theta\,(\Delta I_{ct} \times \text{WRLURI}_c) + \delta\,\text{WRLURI}_c + \alpha X_{ct_0} + \gamma_{dt} + \varepsilon_{ct}, \tag{11}$$

where WRLURI_c measures the restrictiveness of land-use regulation in commuting zone c. The specification without controls omits αX_{ct_0}. I instrument ΔI_{ct} and its interaction with WRLURI jointly with the shift–share instrument Δz_{ct} and its interaction with WRLURI, separately for all, college-educated, and non-college immigrants. WRLURI is the 2018 index of Gyourko et al. (2021), in which higher values indicate more restrictive regulation. I assign the index to counties and aggregate counties to commuting zones using population weights.

The coefficient β measures the effect of immigration when WRLURI equals zero, and θ measures how this effect changes with a one-unit increase in regulation. The effect of immigration at a given level of regulation is

$$\frac{\partial\,\Delta y_{ct}}{\partial\,\Delta I_{ct}} = \beta + \theta\,\text{WRLURI}_c. \tag{12}$$

When housing supply is relatively elastic, an increase in demand can be met by new construction, which limits the increase in rents and home values. In more regulated markets, restrictions on development and delays in approvals constrain this adjustment, so the same increase in demand leads to a larger price response. A positive θ is consistent with this mechanism. WRLURI captures the regulatory part of housing supply constraints, but it does not directly measure supply elasticity or geographic limits on development. The interaction therefore shows whether the effect of immigration rises with regulation, which is consistent with evidence that both regulatory and geographic constraints shape housing supply elasticity across markets (Saiz, 2010).

Including WRLURI narrows the sample, because each observation needs both a housing outcome and a regulatory measure. The samples contain 6,328 commuting-zone–year observations for FMR, 6,244 for HPI, 814 for ZORI, and 5,741 for ZHVI, concentrated in metropolitan areas. The estimates therefore describe the areas covered by both datasets rather than all U.S. commuting zones, and comparisons with the main results should account for this change in coverage.

Table A12 reports estimates of θ. For ZHVI and all immigrants, the interaction is positive and statistically significant in all three base years, with and without controls. It ranges from about 1.17 to 1.24 without controls and from 1.69 to 2.11 with controls. In the controlled specifications, a one-unit increase in WRLURI therefore raises the home-value response to a one-percentage-point immigrant inflow by about 1.7 to 2.1 percentage points. For college-educated immigrants, the ZHVI interaction also increases with controls, ranging from 2.51 to 3.22 and remaining significant in all three base years. These

findings are consistent with immigration-related demand raising home values more where regulation limits new construction. They describe how the effect of immigration varies across regulatory environments, not the overall effect of immigration.

The evidence for the other outcomes is less consistent. With 2000 shares, the controlled FMR interactions are positive and significant for all three immigrant groups, and the HPI interaction is positive and significant for college-educated immigrants with 1980 and 2000 shares. These patterns do not hold across all base years. For non-college immigrants, the controlled ZHVI interaction is significant only with 2000 shares. ZORI provides the least precise evidence: none of its interactions is significant once controls are included. The very large negative estimate for non-college immigrants with 2000 shares has an extremely large standard error and should not be read as evidence that regulation reduces the rent response.

Taken together, the results give the strongest support for a relationship between regulation and the effect of immigration on home values, with more limited evidence for rents and the FHFA index. The number of observations does not change when controls are added, so the changes in coefficients do not reflect further sample loss. The supply mechanism nevertheless remains an interpretation rather than a directly estimated change in supply elasticity, and its credibility depends on the strength of the instruments for both immigration and its interaction with WRLURI, especially in the small ZORI sample.

College-educated and non-college inflows estimated jointly. Appendix Tables A13 and A14 repeat the interaction model with the college-educated and non-college inflows, and their interactions with WRLURI, instrumented jointly, as in Section 6.4. These models have four endogenous variables, and their first stage is weak: the joint first-stage F-statistic is between 3.4 and 8.4 for FMR, HPI, and ZHVI without controls (as low as 0.02 for ZORI), and below 4 with controls (about 1 or less with 1990 and 2000 shares). Without controls, the interaction is positive and significant for college-educated inflows in the ZHVI model with 1980 shares (1.64), the ZORI model with 1990 shares (3.35), and the FMR model with 1990 and 2000 shares (1.14 and 1.12), and for non-college inflows in the ZORI model with 1980 and 1990 shares (2.15 and 2.67) and the FMR model with 2000 shares (1.30). With controls, the only significant interactions are for college-educated inflows in the HPI and ZHVI models with 1980 shares (1.61 and 3.06), where they are also significantly larger than the non-college interactions ($p = 0.044$ and 0.048), and for non-college inflows in the FMR model with 2000 shares (1.15). Because the first stage is so weak, these joint interaction estimates cannot attribute the larger home-value response in regulated markets to one education group.

6.9 Native Displacement

If natives leave a commuting zone one-for-one as immigrants arrive, immigration adds no net housing demand, and rents should not respond. If displacement is only partial, native

outflows offset part of the immigrant demand shock, and the housing response is net of this offset (Saiz, 2007). To assess this, I estimate

$$\frac{N_{ct} - N_{c,t-1}}{L_{ct_0}} = \phi \frac{M_{ct} - M_{c,t-1}}{L_{ct_0}} + \gamma_{dt} + e_{ct}, \tag{13}$$

where N_{ct} is the number of natives and M_{ct} the number of immigrants aged 16 to 64 in commuting zone c in year t, and both changes are divided by the base-year working-age population L_{ct_0} (1980, 1990, or 2000). The coefficient ϕ is the change in the number of working-age natives per additional working-age immigrant: $\phi = -1$ means complete displacement and $\phi = 0$ means none. I estimate the model by OLS and by 2SLS, instrumenting the immigrant change with the shift–share prediction. I estimate it for the whole working-age population and separately for the college-educated and non-college populations; in the education-specific models, both the immigrant change and the instrument are built for the education group and compared with natives in the same group. All specifications include division-by-year fixed effects and are weighted by the base-year working-age population, with standard errors clustered by commuting zone. These models do not include the baseline controls.

Both sides of equation (13) are divided by the same fixed, base-year working-age population, unlike the lagged working-age population in equation (1). With a common denominator that is the same in every year, ϕ directly counts natives per immigrant, and the denominator cannot respond to earlier immigration. In equation (1), the lagged working-age population is used instead because it measures the size of the local market just before each year's inflow. Both choices follow Card's approach of dividing changes by a predetermined population, which Peri and Sparber (2011) show does not build in a bias toward finding displacement, unlike specifications in which native employment also enters the measure of immigration.

The OLS estimates in Table A11 are negative and significant for the whole working-age population (-0.22 to -0.35 across base years) and for the non-college population (-0.18 to -0.23). A coefficient of -0.26, as with 1980 shares, implies about one fewer working-age native for every four working-age immigrants who arrive. This suggests partial displacement concentrated among people without college education, far from the one-for-one pattern of -1. The estimates for the college-educated population are small and insignificant. These associations are not causal, because immigrants may settle in places where the number of working-age natives was already falling for demographic or economic reasons.

The 2SLS estimates are less conclusive. For the whole working-age population and the non-college population, the coefficients are negative with 1980 and 1990 shares but positive with 2000 shares, and the college estimates are positive throughout. None is statistically significant. With 1980 shares, the lack of significance reflects lower precision rather than a smaller magnitude: the estimate for the whole working-age population is nearly identical

under OLS (−0.26) and 2SLS (−0.26), but the 2SLS estimate is far less precise. With 1990 shares the 2SLS estimate is smaller than OLS, and with 2000 shares it has the opposite sign. The 2SLS results therefore do not establish that immigration reduces the working-age native population, but they do not rule out economically meaningful reductions.

These estimates measure changes in the local working-age native population, not native out-migration directly. The number of working-age natives also changes as people age into and out of the 16–64 range and through deaths, and education-specific counts change when people move between education groups without leaving their commuting zone. The results therefore give no clear 2SLS evidence that native departures offset immigrant arrivals, but they cannot show that natives do not move in response to immigration.

6.10 Rotemberg Weights

To identify which historical settlement patterns drive the estimates, Table A9 reports Rotemberg weights for the instruments built from 1980, 1990, and 2000 shares. Following Goldsmith-Pinkham et al. (2020), the 2SLS estimate can be written as a weighted sum of the estimates obtained from each origin's share used as a separate instrument:

$$\widehat{\beta}_{SSIV} = \sum_b \omega_b \widehat{\beta}_b, \qquad \sum_b \omega_b = 1.$$

Here ω_b is the weight on origin b, and $\omega_b \widehat{\beta}_b$ is its contribution to the overall estimate. The weights show which parts of the instrument matter most for identification, and therefore where a violation of the identifying assumption would matter most. They are not immigrant population shares. The weights depend only on the first stage (the instrument, the immigration measure, the controls, and the sample), not on the outcome, while the origin-specific estimates $\widehat{\beta}_b$ depend on the outcome. The estimates in Table A9 use the FHFA HPI as the outcome and the main specification for all immigrants with baseline controls.

The weights show heavy reliance on Cuban and Mexican settlement patterns, especially with 1980 shares. Cuba receives a weight of 0.485 and Mexico 0.181, a combined weight of 0.666. Because the 1980 Cuban share is concentrated mainly in one commuting zone (where it reaches 16.6 percent of the working-age population), much of the identifying variation with 1980 shares likely comes from a single market, Miami. As discussed in Section 5.2, the 1980 census predates the Mariel boatlift, so this Cuban share reflects the earlier exile waves. The combined weight of Cuba and Mexico falls to 0.406 with 1990 shares and 0.301 with 2000 shares, although both remain the top two origins. This decline comes mainly from Cuba, whose weight falls to 0.277 and then 0.165. Meanwhile, India's weight rises from 0.021 with 1980 shares to 0.067 with 2000 shares, and the Dominican Republic becomes another important source of identifying variation. Changing the base year therefore changes which historical settlement patterns drive the estimates.

A large weight need not imply a large contribution to the estimated effect. With 1980 shares, Cuba receives the largest weight but contributes only 0.107 to the overall estimate, while Mexico contributes 0.328 because its origin-specific estimate is much larger. The Philippines and India make negative contributions in all three base years despite positive weights. These negative contributions come from negative origin-specific estimates, not negative weights, and they should not be read as the effect of immigration from those countries alone.

These patterns help interpret how the results change with the base year. The positive all-immigrant estimates for HPI, ZORI, and ZHVI persist across base years, while the college estimates vary more across base years. Changing the base year changes which historical exposures receive the most weight, which can change the overall estimate when housing responses differ across places. This table alone, however, cannot show that such reweighting explains the education-specific results or the differences across housing measures; that would require decompositions for each model's immigration measure, sample, and controls.

Under the share-based identifying assumption, the prominence of Cuba and Mexico makes their historical settlement patterns especially important to scrutinize. The smaller reliance on these two origins with later shares tells us how identification is distributed across origins, but does not by itself show that the later shares are more valid.

7 Discussion

The results are consistent with immigration raising housing demand faster than local supply adjusts in the markets most exposed to national inflows. Three features of the evidence support this interpretation. First, the largest estimates are for ZORI and ZHVI, which measure current asking rents and the value of the whole housing stock. Second, estimated one group at a time, the effects come mainly from non-college immigrants, and the rent effect of non-college inflows remains when both groups are estimated jointly. Third, the positive all-immigrant estimates for HPI, ZORI, and ZHVI persist across all three base years, and the home-value response is larger where land-use regulation is stricter.

The data do not identify a single mechanism. Non-college immigrants may rent for longer, form households differently, or settle in markets with fewer vacancies and less elastic supply. They may also work disproportionately in construction, which would offset some of the price pressure (Monras, 2020), and natives may move away or switch between renting and owning. Each estimate is the net equilibrium response after all these adjustments. Identifying the mechanism would require direct measures of rents by market segment, vacancies, building permits, housing completions, crowding, tenure transitions, and native mobility. The education estimates should therefore be read as differences in the short-run local equilibrium response, not as structural housing demand elasticities.

Measurement matters for the conclusions. Annual changes in FMR and ZORI are almost uncorrelated, even though their levels are highly correlated, so the two rent series do not record the same short-run changes. FMRs are smoothed and lagged by construction and target the lower-middle part of the market, which can mute their annual response. The larger ZORI estimates may reflect faster adjustment of market rents, denser markets, or both. The common-sample results show that geography matters: the HPI and ZHVI estimates roughly double when all outcomes are restricted to Zillow-covered commuting zones. Coefficients estimated on different samples should therefore not be compared as if they described the same markets.

The comparison across education groups also calls for caution. Estimated separately, the non-college estimates are larger and more precise than the college estimates, but the college first stage is weak with 1980 and 1990 shares. The joint estimates in Section 6.4 sharpen this picture. Without controls, non-college inflows raise home values significantly more than college-educated inflows, while both groups raise market rents by similar amounts. With controls, the joint first stage is weak, and the equality test rejects only for home values with 1980 shares, where the college-educated effect is larger. The evidence is therefore strongest for a positive rent effect of non-college inflows, and insufficient to establish that the two groups' effects differ once baseline characteristics are controlled for.

The identification evidence supports the design with 2000 shares more than with older shares. With 2000 shares, the instrument does not predict rent growth in earlier decades, and the first stage is strongest. With 1980 shares, exposure predicts slower rent growth in the 1980s and 1990s, and the Rotemberg weights are concentrated on Cuba and Mexico. In all base years, the historical shares are correlated with several baseline characteristics, which justifies the controls, and some origin-specific estimates differ in sign. On the other hand, the reduced form links predicted exposure directly to housing outcomes without depending on the strength of the first stage (Table A4). These facts limit the claims to settings in which the historical shares can plausibly be treated as unrelated to unobserved housing trends.

Several further limitations remain. Annual differences amplify measurement error and may miss slower quantity responses. Because the main origin, Mexico, sends immigrants to the same places year after year, the estimates may combine the short-run response to current inflows with continuing adjustment to past inflows (Jaeger et al., 2018). Many outcomes and specifications increase the chance of isolated significant estimates, so the most persuasive evidence is the repetition of the same sign and pattern across specifications rather than any single p-value. ZORI's limited coverage restricts national generalization, and weighting by the working-age population makes the estimates more representative of people in large markets than of the typical commuting zone. Finally, the study period includes the housing bust and recovery; division-by-year fixed effects remove broad regional cycles but may not absorb every local shock correlated with immigrant destinations.

Several extensions would strengthen the analysis. First, because the 1980-share estimates rely heavily on Cuban and Mexican settlement, re-estimating the model after dropping the Miami commuting zone and, separately, the Cuban component of the instrument would show how much the results depend on these exposures, as Goldsmith-Pinkham et al. (2020) recommend for origins with large Rotemberg weights. Reporting the total of the negative Rotemberg weights and plotting each origin's estimate against the strength of its first stage would complete the decomposition. Second, constructing the national shocks with each commuting zone's own immigrants left out would ensure that a zone's own inflows do not enter its instrument. Third, Anderson–Rubin confidence intervals (Anderson and Rubin, 1949) would provide inference for the college-educated estimates that remains valid when the first stage is weak. Fourth, several robustness checks remain. ACS samples from 2005 to 2011 identify Public Use Microdata Areas (PUMAs) using 2000 boundaries, while later samples use 2010 boundaries, so the crosswalk to commuting zones changes in 2012; estimates that drop the 2011–2012 change would rule out a mechanical effect of this switch. Unweighted estimates would show whether the effects differ across places, since Solon et al. (2015) note that a large gap between weighted and unweighted estimates signals heterogeneous effects, and standard errors clustered by state would allow for correlation across nearby commuting zones with similar origin mixes. Fifth, the pre-trend test could be repeated with non-overlapping windows, only periods that precede each base year, and census median gross rent as the earlier-period outcome. Finally, the land-use regulation model could report the effect of immigration at low and high levels of regulation, together with first-stage diagnostics for both endogenous variables, and the native displacement model could add the baseline controls.

The policy implications are correspondingly conditional. Immigration can create meaningful short-run pressure on local rental and owner-occupied markets when established networks direct new arrivals to places where supply does not expand quickly. This does not imply that restricting immigration is the only or the preferred response. If the price effect reflects inelastic supply, reforms that speed up permitting and construction can absorb the added demand while preserving the labor market and demographic benefits of immigration. The results speak most directly to the need for housing policy to anticipate population growth that is concentrated in particular places.

8 Conclusion

This paper estimates the effect of immigration on housing markets in U.S. commuting zones from 2005 to 2019. I instrument immigrant inflows with a shift–share instrument that combines historical settlement patterns by country of origin with the national growth of each origin group, and I evaluate the instrument with pre-trend tests, balance regressions, three alternative base years, first stages in the full and Zillow samples, reduced forms, joint estimation of the two education groups, and Rotemberg weights.

Immigration raises market rents and home values in the short run. Across the three base years, an immigrant inflow equal to one percent of the working-age population raises Zillow home values by 3.3 to 3.7 percent, Zillow rents by 1.4 to 3.6 percent, and the FHFA house price index by about 1.5 percent, while Fair Market Rents do not respond detectably. Estimated one group at a time, the response comes mainly from immigrants without college education, whose inflows raise Zillow rents by 4.4 to 7.0 percent and Zillow home values by 4.9 to 5.5 percent, while the college-educated estimates are imprecise. When the two groups are estimated jointly with baseline controls, non-college inflows still raise Zillow rents by 5 to 6 percent in every base year, but the joint first stage is weak and equal effects for the two groups cannot be rejected in most specifications. The 2000 shares give the strongest first stage and do not predict rent growth in earlier decades, while identification with older shares relies more heavily on Cuban and Mexican settlement.

The measured effect depends on which part of the housing market a series covers and where it is observed. FMRs, which are lagged and smoothed by construction, show little short-run response, while the estimates for market rents and home values in the dense markets covered by Zillow are larger and statistically significant. When every outcome is restricted to Zillow-covered commuting zones, the HPI and ZHVI estimates roughly double. The paper therefore does not estimate a single national housing elasticity; it estimates weighted local responses for the markets that the instrument reaches.

Immigration places detectable short-run pressure on market rents and home values in historically exposed commuting zones, especially through non-college inflows and where regulation limits new construction. This conclusion holds across several design choices, but it depends on the assumption that historical settlement shares are unrelated to unobserved housing trends, which is most credible for 2000 shares. Future work that directly measures construction, vacancies, tenure, crowding, and native mobility would help show whether the price response comes mainly from slow supply, segmentation of the rental market, household formation, or residential reallocation.

Appendix

The appendix follows the order of the paper. Figure A1 shows the commuting zones covered by ZORI, and Figure A2 shows how much more concentrated immigrants are than natives across commuting zones. Tables A1 and A2 relate the historical origin shares to baseline characteristics, which is why the main specifications include year interactions with the base-year college, manufacturing, construction, and total immigrant share variables.

The next set of exhibits connects measurement to identification. Table A3 compares the housing outcomes in levels and changes. Table A4 reports the reduced-form effect of the instruments on housing outcomes. Table A5 checks instrument strength in the Zillow rent sample. Figures A3 through A6 show the common housing cycle, the noisy annual changes, and the persistent cross-sectional relationships discussed in Section 6.1.

The final tables cover robustness and scope. Table A6 reports the 2SLS estimates without baseline controls, Tables A7 and A8 estimate the college-educated and non-college inflows jointly, Table A9 reports the origins with the largest Rotemberg weights, Table A10 estimates all four outcomes on the same Zillow-covered observations, Table A11 reports the native displacement estimates, Table A12 reports the WRLURI interaction estimates, and Tables A13 and A14 repeat the interaction model with both education groups instrumented jointly.

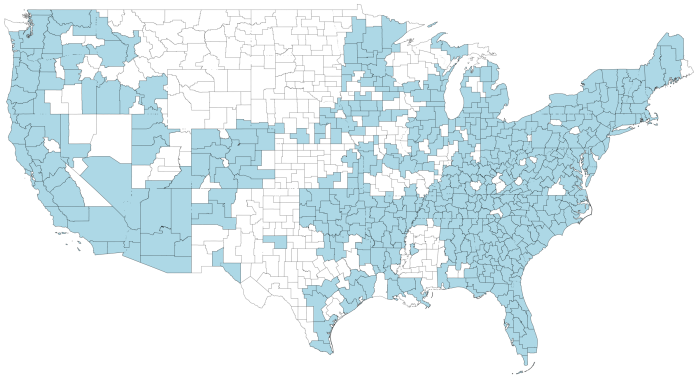


Figure A1: Commuting Zones Covered by ZORI

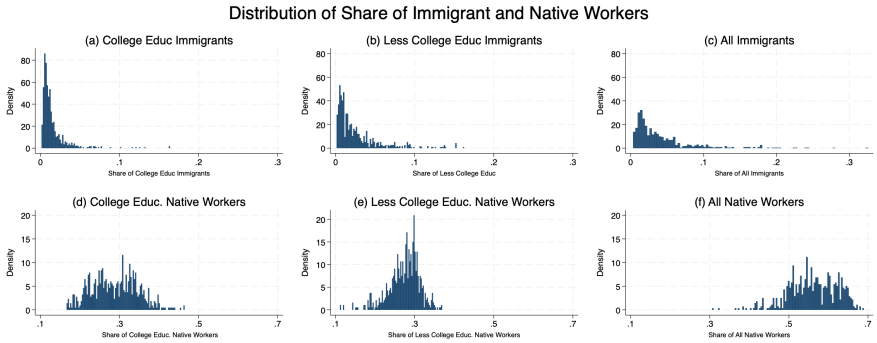


Figure A2: Distribution of Immigrant and Native Shares of the Working-Age Population Across Commuting Zones

	Mexico	Canada	Cuba	Germany	UK	Philippines	India	China	USSR	Japan
College Share 1980	-0.311	0.168	0.138	0.0854***	0.0637***	0.0353	0.0255	0.0470	-0.0334**	0.0349***
	(0.318)	(0.108)	(0.238)	(0.0306)	(0.0193)	(0.0469)	(0.0161)	(0.0431)	(0.0138)	(0.0106)
Manufacturing Share 1980	-0.424*	-0.134*	-0.0472	0.0217	-0.0196	0.0977*	-0.00651	0.0943	0.000441	0.0141**
	(0.242)	(0.0714)	(0.145)	(0.0187)	(0.0145)	(0.0556)	(0.0100)	(0.0675)	(0.0111)	(0.00639)
Construction Share 1980	1.533	-0.270	-0.117	-0.392***	-0.243***	-0.361*	-0.0706	-0.364*	0.120**	-0.192***
	(1.006)	(0.502)	(0.633)	(0.129)	(0.0818)	(0.217)	(0.0562)	(0.199)	(0.0506)	(0.0432)
Sum of Shares 1980	0.279***	0.0113**	0.231	0.0145***	0.00989***	0.0107	0.00560**	0.0201	0.0156***	0.00539**
	(0.0960)	(0.00533)	(0.141)	(0.00455)	(0.00218)	(0.00806)	(0.00272)	(0.0124)	(0.00348)	(0.00225)
Observations	722	722	722	722	722	722	722	722	722	722

Notes: Each column reports a separate cross-sectional regression, across the 722 commuting zones, of one origin's 1980 share on the 1980 characteristics listed. Robust standard errors are reported in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table A1: Historical Origin Shares and Baseline Characteristics, 1980

Panel A: Immigrant Country Shares, 1990										
	Mexico	Germany	Canada	United Kingdom	Philippines	China	Vietnam	Korea	India	Cuba
College Share 1990	0.132	0.0352	0.0455	0.0786***	0.0230	-0.00401	0.234**	0.0376	0.0629	0.499
	(0.507)	(0.0257)	(0.0448)	(0.0193)	(0.0813)	(0.0874)	(0.0971)	(0.0511)	(0.0395)	(0.430)
Manufacturing Share 1990	-0.991***	-0.0273**	-0.107***	-0.0436***	0.131	0.200**	-0.0402	0.000192	-0.00884	0.0650
	(0.280)	(0.0122)	(0.0314)	(0.0113)	(0.0862)	(0.100)	(0.0350)	(0.0264)	(0.0271)	(0.157)
Construction Share 1990	1.510	-0.112	0.0801	-0.247***	-0.468	-0.461	-0.971**	-0.215	-0.248*	-1.847
	(1.880)	(0.104)	(0.202)	(0.0722)	(0.316)	(0.303)	(0.410)	(0.152)	(0.147)	(1.481)
Sum of Shares 1990	0.487***	0.00320	0.00795*	0.00734***	0.0172	0.0230	0.00846	0.0222	0.00703	0.249
	(0.144)	(0.00307)	(0.00473)	(0.00161)	(0.0164)	(0.0193)	(0.00745)	(0.0138)	(0.00633)	(0.155)
Observations	722	722	722	722	722	722	722	722	722	722
Panel B: Immigrant Country Shares, 2000										
	Mexico	Canada	Germany	Philippines	China	United Kingdom	India	Vietnam	Korea	El Salvador
College Share 2000	-0.411	0.105***	0.0305**	0.0271	-0.00797	0.0809***	0.125	0.355**	0.126*	0.330
	(0.558)	(0.0360)	(0.0148)	(0.103)	(0.131)	(0.0187)	(0.0885)	(0.151)	(0.0744)	(0.227)
Manufacturing Share 2000	-0.463	-0.119***	-0.0368***	0.208*	0.303**	-0.0594***	0.00854	-0.0610	-0.0270	-0.0653
	(0.405)	(0.0271)	(0.0110)	(0.112)	(0.137)	(0.0133)	(0.0738)	(0.0868)	(0.0468)	(0.106)
Construction Share 2000	1.346	-0.165	-0.0722	-0.660*	-0.616	-0.232***	-0.600	-1.667**	-0.558**	-1.284*
	(2.549)	(0.165)	(0.0658)	(0.354)	(0.498)	(0.0747)	(0.437)	(0.764)	(0.245)	(0.696)
Sum of Shares 2000	0.491***	0.00538	0.00220	0.0277	0.0390	0.00593***	0.00698	0.00982	0.0284*	0.0538**
	(0.152)	(0.00403)	(0.00196)	(0.0187)	(0.0241)	(0.00160)	(0.0102)	(0.0116)	(0.0162)	(0.0247)
Observations	722	722	722	722	722	722	722	722	722	722

Notes: Each column reports a separate cross-sectional regression, across the 722 commuting zones, of one origin's base-year share on the base-year characteristics listed. Robust standard errors are reported in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table A2: Historical Origin Shares and Baseline Characteristics, 1990 and 2000

	Log Changes				Levels			
	Log Differences, weighted				Log of Outcomes, weighted			
	HPI	FMR Rents	Zillow Rents	Zillow Home Values	HPI	FMR Rents	Zillow Rents	Zillow Home Values
HPI	1				1			
FMR Rents	0.1223	1			0.7056	1		
Zillow Rents	0.2876	0.0679	1		0.6887	0.9157	1	
Zillow Home Values	0.4314	-0.1053	0.2224	1	0.7353	0.9437	0.9351	1
	Log Differences, unweighted				Log Levels, unweighted			
	HPI	FMR Rents	Zillow Rents	Zillow Home Values	HPI	FMR Rents	Zillow Rents	Zillow Home Values
HPI	1				1			
FMR Rents	0.0519	1			0.5882	1		
Zillow Rents	0.3663	0.0451	1		0.5867	0.8214	1	
Zillow Home Values	0.5682	-0.0391	0.2745	1	0.6362	0.8710	0.8322	1

Notes: The table reports pairwise correlation coefficients for housing market outcomes. Log changes are first differences in logs. Weighted correlations use the 1980 working-age population as weights.

Table A3: Correlation Coefficients in Levels and Changes

	FMR	FHFA HPI	ZORI	ZHVI
Panel A: All Immigrants				
Immigration instrument	0.4972	2.6609***	4.2795***	6.5641***
	(1.0034)	(1.0156)	(1.2869)	(0.9608)
Panel B: College-Educated Immigrants				
Immigration instrument	-0.4141	2.7502***	-2.0153***	3.7322
	(1.1584)	(0.9223)	(0.6089)	(2.6511)
Panel C: Non-College Immigrants				
Immigration instrument	0.5873	0.2384	6.2329***	2.7371***
	(0.6296)	(0.5847)	(0.9067)	(0.4757)
Observations	10,108	9,126	868	8,095

Notes: All regressions absorb division×year fixed effects and use analytic weights based on the 1980 working-age population. Controls include the base-year college share, manufacturing share, construction share, and sum of origin shares, each interacted with year indicators. Standard errors clustered by commuting zone are reported in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$. The instrument uses 1980 shares and the national shock for the indicated group.

Table A4: Reduced-Form Regressions on the Shift–Share Instrument, 1980 Shares

	Fixed Shares	
	FE	FE + Controls
Panel A: 1980 Shares		
All Immigrants		
Predicted immigration (Δz_{ct})	1.7990	3.0459
	(0.3634)	(0.5923)
F-statistic	25	26
Observations	868	868
College-Educated Immigrants		
Predicted immigration (Δz_{ct})	0.7666	0.9395
	(0.2932)	(0.3119)
F-statistic	7	9
Observations	868	868
Non-College Immigrants		
Predicted immigration (Δz_{ct})	1.2640	1.4214
	(0.3037)	(0.2112)
F-statistic	17	45
Observations	868	868
Panel B: 1990 Shares		
All Immigrants		
Predicted immigration (Δz_{ct})	1.2287	2.2923
	(0.3911)	(0.5879)
F-statistic	10	15
Observations	868	868
College-Educated Immigrants		
Predicted immigration (Δz_{ct})	0.5329	0.7660
	(0.3165)	(0.4443)
F-statistic	3	3
Observations	868	868
Non-College Immigrants		
Predicted immigration (Δz_{ct})	0.8775	1.1218
	(0.1931)	(0.2088)
F-statistic	21	29
Observations	868	868
Panel C: 2000 Shares		
All Immigrants		
Predicted immigration (Δz_{ct})	0.6970	1.3558
	(0.2523)	(0.3339)
F-statistic	8	16
Observations	868	868
College-Educated Immigrants		
Predicted immigration (Δz_{ct})	0.4131	0.9526
	(0.1559)	(0.2601)
F-statistic	7	13
Observations	868	868
Non-College Immigrants		
Predicted immigration (Δz_{ct})	0.4233	0.6194
	(0.1110)	(0.1439)
F-statistic	15	19
Observations	868	868

Notes: The table reports first-stage estimates in the Zillow rent sample using fixed 1980, 1990, and 2000 shares. Controls include the base-year college share, manufacturing share, construction share, and sum of origin shares, each interacted with year indicators. Standard errors clustered by commuting zone are reported in parentheses. The F-statistic is the squared ratio of the first-stage coefficient to its clustered standard error.

Table A5: First-Stage Estimates in the Zillow Rent Sample Using 1980, 1990, and 2000 Shares

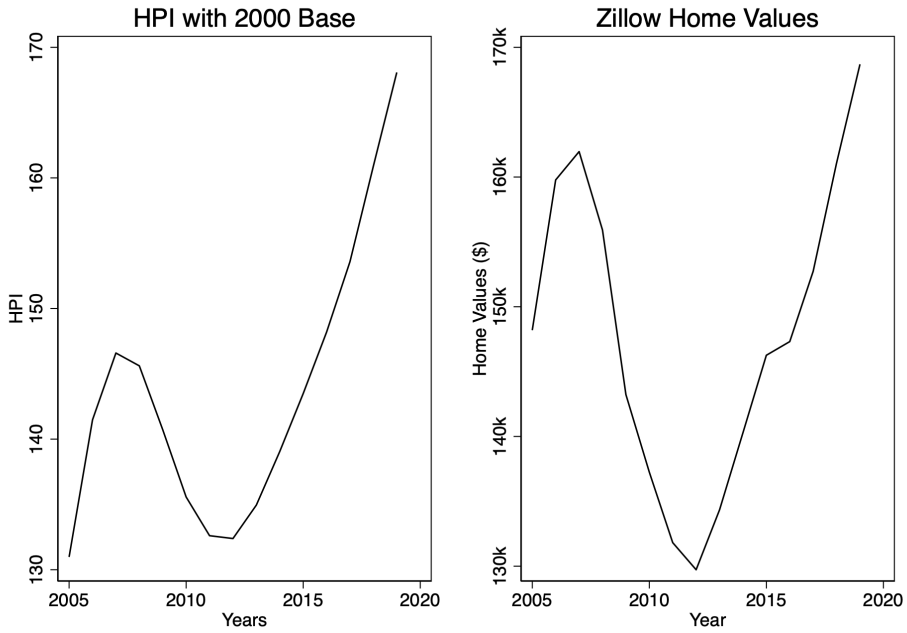


Figure A3: Trends in FHFA HPI and Zillow Home Values

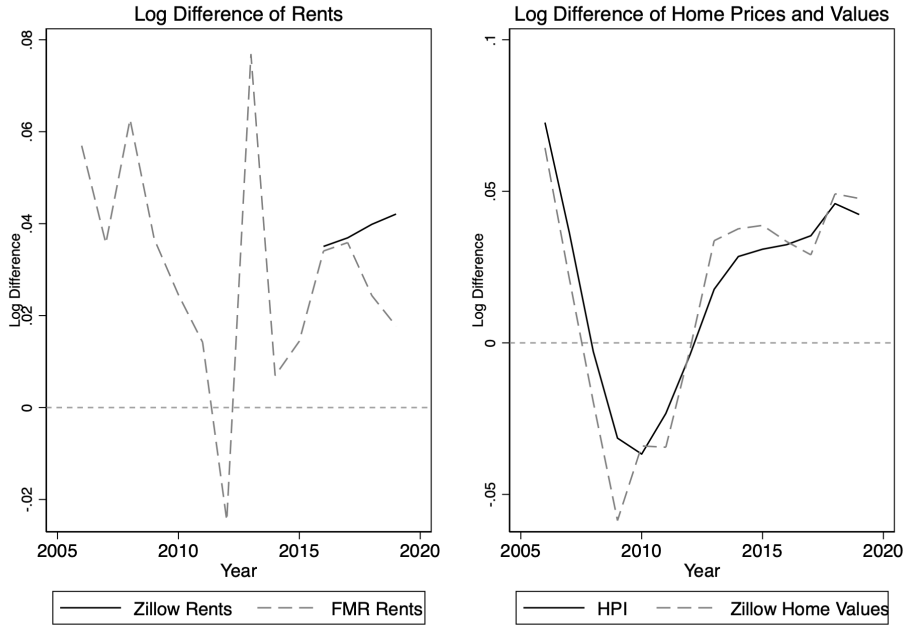


Figure A4: Annual Log Changes in Rents, House Prices, and Home Values

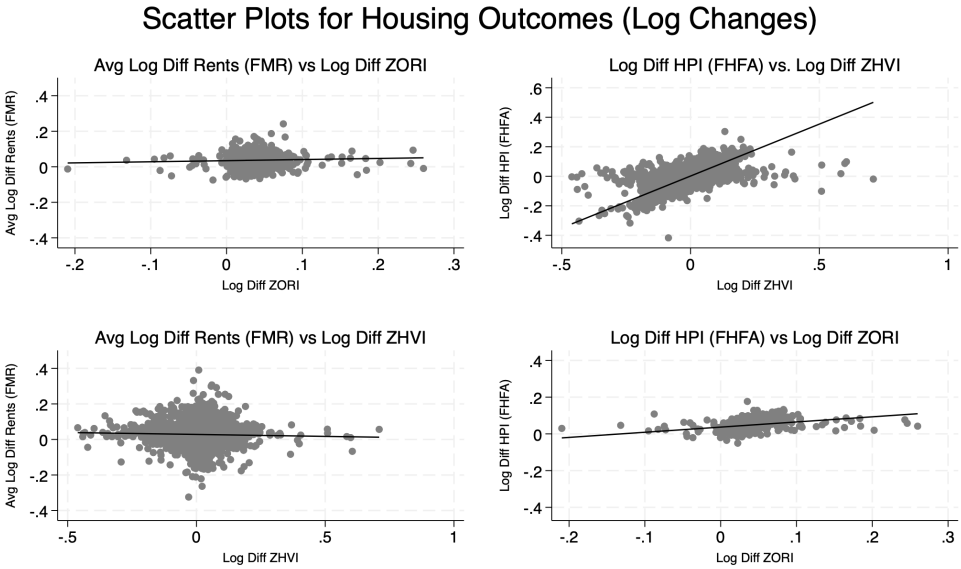


Figure A5: Scatter Plots for Housing Outcomes in Log Changes

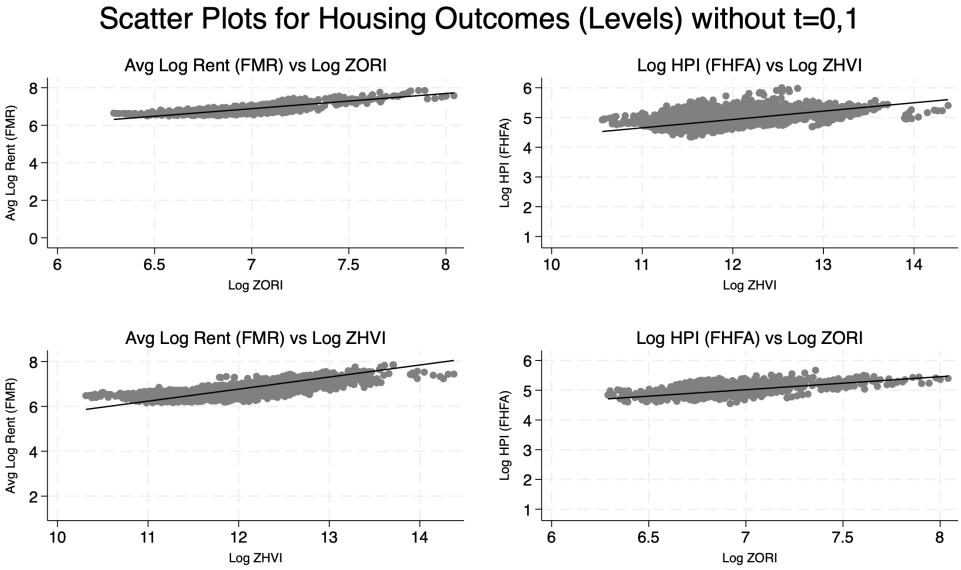


Figure A6: Scatter Plots for Housing Outcomes in Levels Excluding $t = 0, 1$

	FMR (1)	FHFA HPI (2)	ZORI (3)	ZHVI (4)
Panel A: 1980 Shares				
All immigrants	0.0989	0.6313	2.4313***	1.9506*
	(0.6904)	(0.5704)	(0.4985)	(1.1667)
College-educated immigrants	0.3204	-0.7008	0.2564	-0.1658
	(0.8112)	(0.7019)	(0.5623)	(1.7708)
Non-college immigrants	-0.0644	2.5454**	5.0903**	5.2293**
	(0.8414)	(1.1997)	(1.9918)	(2.1038)
Panel B: 1990 Shares				
All immigrants	0.3877	0.1628	4.0210***	1.1468
	(0.7301)	(0.6479)	(1.0541)	(1.2302)
College-educated immigrants	1.0352	-1.5769	2.4566**	-1.8819
	(0.8880)	(0.9964)	(1.0849)	(2.2059)
Non-college immigrants	0.2627	2.5353*	7.2096***	5.9566**
	(1.3084)	(1.3803)	(2.7001)	(2.5498)
Panel C: 2000 Shares				
All immigrants	0.0371	-0.1795	4.6129***	0.7101
	(0.6413)	(0.5183)	(1.4439)	(0.9788)
College-educated immigrants	0.6322	-1.6878**	3.1162*	-1.6611
	(0.7328)	(0.7732)	(1.5908)	(1.4758)
Non-college immigrants	-0.5449	2.1700*	9.2247**	4.8163**
	(1.3759)	(1.1707)	(3.5715)	(2.3163)
Observations	10,108	9,126	868	8,095

Notes: The table reports 2SLS estimates of equation (1) without baseline controls. All regressions include division×year fixed effects and use analytic weights based on the 1980 working-age population. Standard errors clustered by commuting zone are reported in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table A6: 2SLS Estimates Without Baseline Controls

	FMR (1)	FHFA HPI (2)	ZORI (3)	ZHVI (4)
Panel A: 1980 Shares				
College-educated inflow	0.2265	0.1554	3.0544***	1.3038
	(0.7983)	(0.5787)	(0.7508)	(1.6378)
Non-college inflow	-0.2632	2.4102**	3.4420**	4.0832***
	(0.4836)	(1.0110)	(1.7183)	(0.8740)
p-value: equal inflow coefficients	0.593	0.031	0.862	0.119
F-statistic	8.38	8.29	15.32	8.67
Panel B: 1990 Shares				
College-educated inflow	0.8246	-0.6370	5.1761***	-0.1121
	(0.8910)	(0.7215)	(1.0740)	(1.7954)
Non-college inflow	-0.7379	3.3045***	2.7907	6.0934***
	(0.6722)	(1.0455)	(2.1789)	(1.2140)
p-value: equal inflow coefficients	0.103	0.002	0.364	0.003
F-statistic	16.30	16.19	7.08	16.63
Panel C: 2000 Shares				
College-educated inflow	0.5050	-1.2815*	5.8490***	-0.9792
	(0.7541)	(0.6543)	(1.5112)	(1.3934)
Non-college inflow	-1.1605	3.7271***	3.7210*	6.0160***
	(0.7950)	(0.8880)	(2.1181)	(1.3401)
p-value: equal inflow coefficients	0.051	0.000	0.372	0.000
F-statistic	29.48	29.33	12.57	29.11
Observations	10,108	9,126	868	8,095

Notes: Each column reports one 2SLS regression in which the college-educated and non-college net inflows are both endogenous and are instrumented jointly with the college-educated and non-college shift–share instruments for the indicated base year. The p-value tests whether the two inflow coefficients are equal. The F-statistic is the Kleibergen–Paap rk Wald F-statistic (Kleibergen and Paap, 2006) for the joint first stage. All regressions include division×year fixed effects and use analytic weights based on the base-year working-age population (1980, 1990, or 2000). Standard errors clustered by commuting zone are reported in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table A7: Joint 2SLS Estimates for College-Educated and Non-College Inflows Without Baseline Controls

	FMR (1)	FHFA HPI (2)	ZORI (3)	ZHVI (4)
Panel A: 1980 Shares				
College-educated inflow	0.0865	2.6456***	-1.6892	5.4262***
	(1.3500)	(0.9592)	(2.3989)	(1.7128)
Non-college inflow	0.9516	-2.4896*	5.3468***	-1.0565
	(1.3342)	(1.5042)	(2.0268)	(1.7025)
p-value: equal inflow coefficients	0.728	0.010	0.102	0.042
F-statistic	8.04	8.06	2.84	7.81
Panel B: 1990 Shares				
College-educated inflow	1.7435	1.2721	-1.3927	-0.7933
	(2.8446)	(2.1417)	(3.8731)	(3.7163)
Non-college inflow	-0.4208	-0.4264	6.1545**	6.0515
	(2.9895)	(2.8079)	(2.6886)	(3.7361)
p-value: equal inflow coefficients	0.705	0.725	0.242	0.341
F-statistic	1.44	1.41	1.14	1.29
Panel C: 2000 Shares				
College-educated inflow	-0.5598	-0.3320	1.0320	-2.0181
	(1.5931)	(1.2903)	(2.1598)	(2.0480)
Non-college inflow	0.4211	1.4994	5.7804***	5.6506*
	(1.9348)	(1.5172)	(1.5289)	(3.1327)
p-value: equal inflow coefficients	0.762	0.501	0.161	0.126
F-statistic	2.11	2.09	4.08	2.01
Observations	10,108	9,126	868	8,095

Notes: Each column reports one 2SLS regression in which the college-educated and non-college net inflows are both endogenous and are instrumented jointly with the college-educated and non-college shift–share instruments for the indicated base year. The p-value tests whether the two inflow coefficients are equal. The F-statistic is the Kleibergen–Paap rk Wald F-statistic (Kleibergen and Paap, 2006) for the joint first stage. All regressions include division×year fixed effects and the baseline controls (the base-year college share, manufacturing share, construction share, and sum of origin shares, each interacted with year indicators), and use analytic weights based on the base-year working-age population (1980, 1990, or 2000). Standard errors clustered by commuting zone are reported in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table A8: Joint 2SLS Estimates for College-Educated and Non-College Inflows With Baseline Controls

Origin	Rank	ω	$\lvert\omega\rvert$	β_{origin}	$\beta_{\text{contribution}}$
Panel A: 1980 Shares					
Cuba	1	0.485	0.485	0.220	0.107
Mexico	2	0.181	0.181	1.810	0.328
Philippines	3	0.030	0.030	-4.062	-0.122
Trinidad and Tobago	4	0.028	0.028	0.198	0.005
Poland	5	0.025	0.025	3.827	0.094
Greece	6	0.024	0.024	2.457	0.059
Colombia	7	0.023	0.023	5.293	0.119
Italy	8	0.021	0.021	3.671	0.078
India	9	0.021	0.021	-1.197	-0.025
Ireland	10	0.019	0.019	1.195	0.023
Panel B: 1990 Shares					
Cuba	1	0.277	0.277	0.233	0.064
Mexico	2	0.129	0.129	1.189	0.153
Dominican Republic	3	0.050	0.050	3.561	0.179
Colombia	4	0.043	0.043	1.666	0.071
India	5	0.042	0.042	-1.192	-0.050
Portugal	6	0.038	0.038	-1.889	-0.071
Jamaica	7	0.035	0.035	0.538	0.019
Nicaragua	8	0.033	0.033	-6.937	-0.230
China	9	0.029	0.029	2.317	0.068
Philippines	10	0.028	0.028	-5.815	-0.165
Panel C: 2000 Shares					
Cuba	1	0.165	0.165	0.240	0.039
Mexico	2	0.136	0.136	0.289	0.039
India	3	0.067	0.067	-1.138	-0.076
Dominican Republic	4	0.047	0.047	3.709	0.175
Colombia	5	0.043	0.043	0.859	0.037
Jamaica	6	0.037	0.037	0.280	0.010
China	7	0.035	0.035	1.368	0.048
Venezuela	8	0.032	0.032	1.193	0.039
Philippines	9	0.028	0.028	-5.872	-0.163
Honduras	10	0.025	0.025	-1.064	-0.027

Notes: The table reports the ten highest-ranked origin countries by Rotemberg weight for each base-year specification. ω is the Rotemberg weight, $|\omega|$ is its absolute value, β_{origin} is the origin-specific estimate, and $\beta_{\text{contribution}}$ is the origin's contribution to the aggregate estimate. The weights are computed from the first stage for all immigrants with baseline controls, and β_{origin} uses the FHFA HPI as the outcome.

Table A9: Top 10 Rotemberg-Weight Origins by Base Year

	FMR	FHFA HPI	ZORI	ZHVI
Panel A: 1980 Shares				
All Immigrants				
Actual inflows, ΔI_{ct}	-0.4529	2.2452**	1.4140***	5.6065***
	(0.5135)	(0.8996)	(0.4022)	(0.8090)
College-Educated Immigrants				
Actual inflows, ΔI_{ct}	-1.9577*	4.5698**	-2.2427**	8.7259***
	(1.0799)	(2.0254)	(1.1168)	(2.3182)
Non-College Immigrants				
Actual inflows, ΔI_{ct}	-0.1272	1.5443	4.2695***	7.4277***
	(0.7062)	(1.0600)	(0.8173)	(1.5214)
Observations	814	814	814	814
Panel B: 1990 Shares				
All Immigrants				
Actual inflows, ΔI_{ct}	0.1616	3.0542**	2.6029***	6.6927***
	(0.8869)	(1.5422)	(0.5677)	(1.0706)
College-Educated Immigrants				
Actual inflows, ΔI_{ct}	-2.5383	6.1032	-1.3791	11.2247**
	(2.2539)	(3.6969)	(1.8905)	(4.8450)
Non-College Immigrants				
Actual inflows, ΔI_{ct}	0.8092	2.9795*	5.6518***	9.8090***
	(1.2964)	(1.6940)	(1.0286)	(1.8934)
Observations	814	814	814	814
Panel C: 2000 Shares				
All Immigrants				
Actual inflows, ΔI_{ct}	0.3965	2.8912*	3.5846***	6.6151***
	(0.9924)	(1.6134)	(0.7740)	(1.1890)
College-Educated Immigrants				
Actual inflows, ΔI_{ct}	-1.3673	4.1171**	1.7267	9.1533***
	(1.9550)	(1.7780)	(1.3399)	(2.7606)
Non-College Immigrants				
Actual inflows, ΔI_{ct}	1.1692	3.9926	6.8408***	10.3177***
	(1.6159)	(2.4908)	(1.5439)	(2.3271)
Observations	814	814	814	814

Notes: The table reports all-controls 2SLS estimates on the common sample for which Zillow rent data are available. All regressions absorb division×year fixed effects and use analytic weights based on the 1980 working-age population. Controls include the base-year college share, manufacturing share, construction share, and sum of origin shares, each interacted with year indicators. Standard errors clustered by commuting zone are reported in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$. The common sample has 814 observations, fewer than the 868 ZORI observations, because it also requires FMR, HPI, and ZHVI.

Table A10: 2SLS Estimates Using the Common Zillow Rent Sample

	Dependent variable: change in working-age natives, ΔN_{ct}					
	OLS			2SLS		
	All	College	Non-College	All	College	Non-College
Panel A: 1980 Shares						
Actual inflows, ΔI_{ct}	-0.2581***	0.0728	-0.2242***	-0.2623	0.2806	-0.4940
	(0.0543)	(0.0633)	(0.0456)	(0.3934)	(0.3216)	(0.3023)
Observations	10,108	10,108	10,108	10,108	10,108	10,108
Panel B: 1990 Shares						
Actual inflows, ΔI_{ct}	-0.2178***	0.0243	-0.1779***	-0.0841	0.1159	-0.0751
	(0.0416)	(0.0462)	(0.0346)	(0.2878)	(0.1951)	(0.2824)
Observations	10,108	10,108	10,108	10,108	10,108	10,108
Panel C: 2000 Shares						
Actual inflows, ΔI_{ct}	-0.3513***	-0.0326	-0.2323***	0.1152	0.3380	0.1121
	(0.0431)	(0.0542)	(0.0392)	(0.3150)	(0.2265)	(0.3343)
Observations	10,108	10,108	10,108	10,108	10,108	10,108

Notes: The dependent variable is the annual change in the number of working-age natives, and the regressor is the change in the number of working-age immigrants, both scaled by the base-year working-age population of the indicated share year. A coefficient of -1 corresponds to complete displacement and 0 to none. In the All columns the regressor is the change in all working-age immigrants; in the College (Non-College) columns, both the immigrant inflow and the instrument are constructed for college-educated (non-college) immigrants and compared with natives in the same education group. The 2SLS columns instrument the immigrant inflow with the shift–share prediction built from baseline immigrant shares. All regressions absorb division×year fixed effects, are weighted by the base-year working-age population, and exclude baseline controls. Standard errors clustered by commuting zone are reported in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table A11: Native Displacement: OLS and 2SLS Estimates Using 1980, 1990, and 2000 Shares

$\Delta I_{ct} \times \text{WRLURI}_c$	FMR	FHFA HPI	ZORI	ZHVI
1980 Shares				
Panel A: Interaction Model				
All immigrants	0.6573	0.4967	2.3524	1.2446***
	(0.6049)	(0.4766)	(1.9348)	(0.4271)
College-educated immigrants	0.9946	0.5353	1.0255	1.5746**
	(0.7582)	(0.5058)	(1.1986)	(0.7047)
Non-college immigrants	0.1842	0.5870	2.3775	0.8475
	(0.5256)	(0.6271)	(1.6903)	(0.9808)
Observations	6,328	6,244	814	5,741
Panel B: Interaction Model with Controls				
All immigrants	0.3881	0.4637	-0.9062	1.6868***
	(0.4605)	(0.3930)	(1.6423)	(0.4959)
College-educated immigrants	0.6743	1.6386*	1.3855	3.2202***
	(0.6577)	(0.9809)	(2.2713)	(0.9293)
Non-college immigrants	0.0246	-0.3884	0.7786	0.6438
	(0.6109)	(0.4655)	(1.3137)	(0.9984)
Observations	6,328	6,244	814	5,741
1990 Shares				
Panel A: Interaction Model				
All immigrants	0.8434	0.3610	5.7585	1.1745***
	(0.6018)	(0.5168)	(5.7704)	(0.4439)
College-educated immigrants	1.3078*	0.2724	1.6381	1.0677
	(0.7681)	(0.6837)	(1.2805)	(0.9126)
Non-college immigrants	0.6275	0.3656	3.6007	1.0124
	(0.5925)	(0.5778)	(2.1969)	(0.9469)
Observations	6,328	6,244	814	5,741
Panel B: Interaction Model with Controls				
All immigrants	0.5806	0.4662	-1.7432	1.8007***
	(0.4947)	(0.4470)	(4.2662)	(0.4808)
College-educated immigrants	1.7963	2.1223	1.0982	3.0230**
	(1.9208)	(1.4430)	(1.9445)	(1.3992)
Non-college immigrants	0.3427	-0.4399	-0.7272	0.6703
	(0.7061)	(0.4997)	(2.2873)	(0.9532)
Observations	6,328	6,244	814	5,741
2000 Shares				
Panel A: Interaction Model				
All immigrants	1.1493**	0.4034	8.4474	1.1789***
	(0.5688)	(0.4548)	(13.3722)	(0.3997)
College-educated immigrants	1.6656**	0.2310	2.6336*	0.8063
	(0.6922)	(0.6258)	(1.4343)	(0.6452)
Non-college immigrants	1.5218**	0.8040	9.9577	1.5599**
	(0.6577)	(0.5013)	(23.9976)	(0.7400)
Observations	6,328	6,244	814	5,741
Panel B: Interaction Model with Controls				
All immigrants	1.0747**	0.7387*	-1.5918	2.1110***
	(0.5078)	(0.4133)	(7.9785)	(0.4381)
College-educated immigrants	2.3995*	1.4282**	1.8504	2.5052**
	(1.2868)	(0.7221)	(1.5332)	(1.0461)
Non-college immigrants	1.4510**	0.2690	-44.8142	1.5780**
	(0.6893)	(0.4182)	(265.3861)	(0.7004)
Observations	6,328	6,244	814	5,741

Notes: The table reports 2SLS estimates of the interaction coefficient θ on $\Delta I_{ct} \times \text{WRLURI}_c$ from equation (11), estimated separately for each immigrant group, outcome, and base year. ΔI_{ct} and its interaction with WRLURI are instrumented jointly using the shift–share instrument Δz_{ct} and its interaction with WRLURI. Panel A reports the interaction model without baseline controls; Panel B adds the baseline controls (the base-year college share, manufacturing share, construction share, and sum of origin shares, each interacted with year indicators). WRLURI is the 2018 index of Gyourko et al. (2021). All regressions absorb division×year fixed effects and use analytic weights based on the 1980 working-age population. Standard errors clustered by commuting zone are reported in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table A12: 2SLS Estimates of the Immigration–WRLURI Interaction Using 1980, 1990, and 2000 Shares

	FMR (1)	FHFA HPI (2)	ZORI (3)	ZHVI (4)
Panel A: 1980 Shares				
College-educated inflow	-0.5360	-0.0604	0.8320	0.2700
	(0.7725)	(0.6084)	(2.0161)	(1.2539)
Non-college inflow	-0.2937	1.9514*	1.8928	3.2698***
	(0.7097)	(1.0630)	(1.3019)	(1.0476)
College-educated inflow × WRLURI	0.9770	0.5590	1.8774	1.6398**
	(0.7016)	(0.4302)	(1.5008)	(0.7146)
Non-college inflow × WRLURI	0.0016	0.5078	2.1532**	0.6676
	(0.5428)	(0.6351)	(1.0092)	(0.7751)
WRLURI	-0.0018	-0.0044**	-0.0039	-0.0076*
	(0.0029)	(0.0019)	(0.0067)	(0.0046)
F-statistic	3.39	3.41	2.70	3.48
Panel B: 1990 Shares				
College-educated inflow	0.0904	-0.7189	1.8407	-0.8342
	(0.9331)	(0.7578)	(1.7048)	(1.5172)
Non-college inflow	-1.1735	3.1406***	0.7496	5.2323***
	(0.9008)	(1.1697)	(1.9726)	(1.4296)
College-educated inflow × WRLURI	1.1435*	0.2189	3.3498**	0.8142
	(0.6839)	(0.6648)	(1.5048)	(0.9928)
Non-college inflow × WRLURI	0.4858	0.2147	2.6725*	0.7729
	(0.7639)	(0.5295)	(1.5357)	(1.0228)
WRLURI	-0.0029	-0.0015	-0.0100*	-0.0019
	(0.0025)	(0.0023)	(0.0057)	(0.0047)
F-statistic	6.07	6.10	2.45	6.39
Panel C: 2000 Shares				
College-educated inflow	-0.1440	-1.4248**	-1.0947	-1.6021
	(0.7110)	(0.6742)	(12.1459)	(1.1581)
Non-college inflow	-2.3018**	3.2227***	-6.0917	4.6805***
	(1.0301)	(1.0724)	(31.9572)	(1.3148)
College-educated inflow × WRLURI	1.1152*	0.1593	8.1163	0.5486
	(0.6200)	(0.6598)	(15.6297)	(0.8117)
Non-college inflow × WRLURI	1.3013*	0.6189	12.1321	1.2983
	(0.7306)	(0.4232)	(39.8598)	(0.8510)
WRLURI	-0.0026	-0.0006	-0.0253	-0.0001
	(0.0019)	(0.0020)	(0.0623)	(0.0034)
F-statistic	7.98	8.01	0.02	8.40
Observations	6,328	6,244	814	5,741

Notes: Each column reports one 2SLS regression in which the college-educated and non-college net inflows and their interactions with WRLURI are endogenous and are instrumented jointly with the two education-specific shift–share instruments and their interactions with WRLURI. The F-statistic is the Kleibergen–Paap rk Wald F-statistic (Kleibergen and Paap, 2006) for the joint first stage. WRLURI is the 2018 index of Gyourko et al. (2021), aggregated from counties to commuting zones with population weights. All regressions include WRLURI and division×year fixed effects, and use analytic weights based on the base-year working-age population (1980, 1990, or 2000). Standard errors clustered by commuting zone are reported in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table A13: Joint 2SLS Estimates of the Immigration–WRLURI Interaction Without Baseline Controls

	FMR (1)	FHFA HPI (2)	ZORI (3)	ZHVI (4)
Panel A: 1980 Shares				
College-educated inflow	-0.5688	1.1438	-1.4261	2.6597*
	(1.2319)	(1.2222)	(2.6147)	(1.4940)
Non-college inflow	0.9035	-2.1257	4.6021**	-1.6028
	(1.2576)	(1.3841)	(1.8981)	(1.7014)
College-educated inflow × WRLURI	0.7999	1.6149*	-0.0868	3.0608***
	(0.6713)	(0.8447)	(1.7373)	(0.9884)
Non-college inflow × WRLURI	-0.2192	-0.7037	0.7732	0.0650
	(0.6336)	(0.6471)	(1.4647)	(0.8375)
WRLURI	-0.0010	-0.0050*	0.0017	-0.0124***
	(0.0021)	(0.0029)	(0.0071)	(0.0033)
p-value: equal inflow coefficients	0.505	0.048	0.111	0.103
p-value: equal interaction coefficients	0.333	0.044	0.699	0.048
F-statistic	3.62	3.66	1.49	3.75
Panel B: 1990 Shares				
College-educated inflow	1.3196	0.6821	1.2941	-1.2275
	(2.8181)	(2.1837)	(2.6264)	(2.8459)
Non-college inflow	-1.5625	-1.0248	5.8353**	4.9327
	(4.0677)	(3.9497)	(2.7699)	(4.6084)
College-educated inflow × WRLURI	1.4231	1.6056	-1.4145	1.2575
	(1.4257)	(1.2867)	(2.3464)	(1.9644)
Non-college inflow × WRLURI	0.3824	-0.5135	-0.5770	0.2292
	(0.7797)	(0.6391)	(1.8933)	(1.0726)
WRLURI	-0.0030	-0.0044	0.0072	-0.0048
	(0.0042)	(0.0036)	(0.0086)	(0.0066)
p-value: equal inflow coefficients	0.670	0.771	0.259	0.397
p-value: equal interaction coefficients	0.526	0.173	0.718	0.622
F-statistic	0.51	0.52	1.05	0.50
Panel C: 2000 Shares				
College-educated inflow	-0.2678	0.0161	10.0654	-1.8715
	(2.1313)	(1.9372)	(29.3320)	(2.1955)
Non-college inflow	-1.0829	0.3350	18.9278	3.9886
	(2.9764)	(3.0262)	(49.3928)	(3.6304)
College-educated inflow × WRLURI	0.8240	1.1246	-4.1696	1.1711
	(1.1209)	(1.1307)	(16.0294)	(1.2857)
Non-college inflow × WRLURI	1.1465*	0.0532	-19.4228	0.9409
	(0.6255)	(0.3565)	(66.0624)	(0.7324)
WRLURI	-0.0011	-0.0026	0.0234	-0.0038
	(0.0029)	(0.0027)	(0.0849)	(0.0039)
p-value: equal inflow coefficients	0.868	0.948	0.680	0.306
p-value: equal interaction coefficients	0.789	0.351	0.764	0.870
F-statistic	0.56	0.57	0.03	0.54
Observations	6,328	6,244	814	5,741

Notes: Each column reports one 2SLS regression in which the college-educated and non-college net inflows and their interactions with WRLURI are endogenous and are instrumented jointly with the two education-specific shift–share instruments and their interactions with WRLURI. The first p-value tests whether the two inflow coefficients are equal, and the second tests whether the two interaction coefficients are equal. The F-statistic is the Kleibergen–Paap rk Wald F-statistic (Kleibergen and Paap, 2006) for the joint first stage. WRLURI is the 2018 index of Gyourko et al. (2021), aggregated from counties to commuting zones with population weights. All regressions include WRLURI, division×year fixed effects, and the baseline controls (the base-year college share, manufacturing share, construction share, and sum of origin shares, each interacted with year indicators), and use analytic weights based on the base-year working-age population (1980, 1990, or 2000). Standard errors clustered by commuting zone are reported in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table A14: Joint 2SLS Estimates of the Immigration–WRLURI Interaction With Baseline Controls

References

Antonio Accetturo, Francesco Manaresi, Sauro Mocetti, and Elisabetta Olivieri. Don't stand so close to me: The urban impact of immigration. *Regional Science and Urban Economics*, 45:45–56, 2014.

Rodrigo Adão, Michal Kolesár, and Eduardo Morales. Shift-share designs: Theory and inference. *Quarterly Journal of Economics*, 134(4):1949–2010, 2019.

Joseph G. Altonji and David Card. The effects of immigration on the labor market outcomes of less-skilled natives. In John M. Abowd and Richard B. Freeman, editors, *Immigration, Trade, and the Labor Market*, pages 201–234. University of Chicago Press, Chicago, 1991.

T. W. Anderson and Herman Rubin. Estimation of the parameters of a single equation in a complete system of stochastic equations. *Annals of Mathematical Statistics*, 20(1): 46–63, 1949.

Amanda Ang, Eunjee Kwon, and Siqi Zheng. Asian immigrants, school quality, and the U.S. housing market. *Journal of Urban Economics*, 153:103860, 2026.

David Autor, David Dorn, and Gordon Hanson. When work disappears: Manufacturing decline and the falling marriage market value of young men. *American Economic Review: Insights*, 1(2):161–178, 2019.

David H. Autor and David Dorn. The growth of low-skill service jobs and the polarization of the US labor market. *American Economic Review*, 103(5):1553–1597, 2013.

Timothy J. Bartik. *Who Benefits from State and Local Economic Development Policies?* W.E. Upjohn Institute for Employment Research, Kalamazoo, MI, 1991.

Alexander N. Bogin, William M. Doerner, and William D. Larson. Local house price dynamics: New indices and stylized facts. *Real Estate Economics*, 47(2):365–398, 2019.

George J. Borjas. Homeownership in the immigrant population. *Journal of Urban Economics*, 52(3):448–476, 2002.

Kirill Borusyak, Peter Hull, and Xavier Jaravel. Quasi-experimental shift-share research designs. *Review of Economic Studies*, 89(1):181–213, 2022.

David Card. The impact of the Mariel boatlift on the Miami labor market. *Industrial and Labor Relations Review*, 43(2):245–257, 1990.

David Card. Immigrant inflows, native outflows, and the local labor market impacts of higher immigration. *Journal of Labor Economics*, 19(1):22–64, 2001.

Congressional Budget Office. The demographic outlook: 2026 to 2056. Report 61879, Congressional Budget Office, 2026. URL https://www.cbo.gov/publication/61879.

Kathrin Degen and Andreas M. Fischer. Immigration and Swiss house prices. *Swiss Journal of Economics and Statistics*, 153(1):15–36, 2017.

Federal Housing Finance Agency. FHFA house price index methodology. https://www.fhfa.gov/data/hpi, 2024. Accessed 2025.

Paul Goldsmith-Pinkham, Isaac Sorkin, and Henry Swift. Bartik instruments: What, when, why, and how. *American Economic Review*, 110(8):2586–2624, 2020.

Libertad Gonzalez and Francesc Ortega. Immigration and housing booms: Evidence from Spain. *Journal of Regional Science*, 53(1):37–59, 2013.

Joseph Gyourko, Jonathan S. Hartley, and Jacob Krimmel. The local residential land use regulatory environment across U.S. housing markets: Evidence from a new Wharton index. *Journal of Urban Economics*, 124:103337, 2021.

David A. Jaeger, Joakim Ruist, and Jan Stuhler. Shift-share instruments and the impact of immigration. NBER Working Paper 24285, National Bureau of Economic Research, 2018.

Frank Kleibergen and Richard Paap. Generalized reduced rank tests using the singular value decomposition. *Journal of Econometrics*, 133(1):97–126, 2006.

Joan Monras. Immigration and wage dynamics: Evidence from the Mexican peso crisis. *Journal of Political Economy*, 128(8):3017–3089, 2020.

José Luis Montiel Olea and Carolin Pflueger. A robust test for weak instruments. *Journal of Business & Economic Statistics*, 31(3):358–369, 2013.

Abeba Mussa, Uwaoma G. Nwaogu, and Susan Pozo. Immigration and housing: A spatial econometric analysis. *Journal of Housing Economics*, 35:13–25, 2017.

Gianmarco I. P. Ottaviano and Giovanni Peri. The economic value of cultural diversity: Evidence from US cities. *Journal of Economic Geography*, 6(1):9–44, 2006.

Gianmarco I. P. Ottaviano and Giovanni Peri. Rethinking the effect of immigration on wages. *Journal of the European Economic Association*, 10(1):152–197, 2012.

Giovanni Peri and Chad Sparber. Assessing inherent model bias: An application to native displacement in response to immigration. *Journal of Urban Economics*, 69(1):82–91, 2011.

Steven Ruggles, Sarah Flood, Matthew Sobek, Daniel Backman, Grace Cooper, Julia A. Rivera Drew, Stephanie Richards, Renae Rogers, Jonathan Schroeder, and Kari C. W. Williams. IPUMS USA: Version 16.0 [dataset]. Minneapolis, MN: IPUMS, 2025.

Filipa Sá. Immigration and house prices in the UK. *Economic Journal*, 125(587):1393–1424, 2015.

Albert Saiz. Room in the kitchen for the melting pot: Immigration and rental prices. *Review of Economics and Statistics*, 85(3):502–521, 2003.

Albert Saiz. Immigration and housing rents in American cities. *Journal of Urban Economics*, 61(2):345–371, 2007.

Albert Saiz. The geographic determinants of housing supply. *Quarterly Journal of Economics*, 125(3):1253–1296, 2010.

Albert Saiz and Susan Wachter. Immigration and the neighborhood. *American Economic Journal: Economic Policy*, 3(2):169–188, 2011.

Jamie Sharpe. Re-evaluating the impact of immigration on the U.S. rental housing market. *Journal of Urban Economics*, 111:14–34, 2019.

Gary Solon, Steven J. Haider, and Jeffrey M. Wooldridge. What are we weighting for? *Journal of Human Resources*, 50(2):301–316, 2015.

Charles M. Tolbert and Molly Sizer. U.S. commuting zones and labor market areas: A 1990 update. ERS Staff Paper 9614, U.S. Department of Agriculture, Economic Research Service, 1996.

U.S. Department of Housing and Urban Development. Fair market rent documentation system. https://www.huduser.gov/portal/datasets/fmr.html, 2024. Accessed 2025.

Arno J. van der Vlist, Daniel Czamanski, and Henk Folmer. Immigration and urban housing market dynamics: The case of Haifa. *Annals of Regional Science*, 47(3):585–598, 2011.

Zillow Research. Zillow home value index (ZHVI): Methodology. https://www.zillow.com/research/methodology-neural-zhvi-32128/, 2024a. Accessed 2025.

Zillow Research. Zillow observed rent index (ZORI): Methodology. https://www.zillow.com/research/methodology-zori-repeat-rent-27092/, 2024b. Accessed 2025.