Journal of Urban Economics 133 (2023) 103383

Contents lists available at ScienceDirect

Journal of Urban Economics

journal homepage: www.elsevier.com/locate/jue

JUE Insight: The effect of new market-rate housing construction on the low-income housing market☆

Evan Mast

W.E. Upjohn Institute for Employment Research, 300 S. Westnedge Avenue, Kalamazoo, MI 49007, United States

ARTICLE INFO

JEL classification:
R31
R21
R23

Keywords:
Housing supply
Housing affordability
Filtering

ABSTRACT

I illustrate how new market-rate construction loosens the market for lower-quality housing through a series of moves. First, I use address history data to identify 52,000 residents of new multifamily buildings in large cities, their previous address, the current residents of those addresses, and so on for six rounds. The sequence quickly reaches units in below-median income neighborhoods, which account for nearly 40 percent of the sixth round, and similar patterns appear for neighborhoods in the bottom quintile of income or percent white. Next, I use a simple simulation model to roughly quantify these migratory connections under a range of assumptions. Constructing a new market-rate building that houses 100 people ultimately leads 45 to 70 people to move out of below-median income neighborhoods, with most of the effect occurring within three years. These results suggest that the migration ripple effects of new housing will affect a wide spectrum of neighborhoods and loosen the low-income housing market.

1. Introduction

There are a variety of approaches to affordable housing policy. Many government interventions, such as constructing affordable units or requiring income-restricted units in new developments, directly add to the stock of affordable housing. Another set of strategies encourage market-rate housing construction through relaxed land-use regulation or subsidies. Because most new construction is expensive, the latter approach only indirectly affects the stock of affordable housing. These indirect "filtering" effects are nuanced and difficult to observe, but they are crucial for understanding the relative merits of the two approaches and formulating optimal housing policy.

A long, largely theoretical literature has considered filtering in a variety of models (Sweeney, 1974; Braid, 1981; Nathanson, 2019). In the long run, new housing can depreciate and gradually become affordable, as shown empirically by Weicher et al. (2010); Rosenthal (2014), and Liu et al. (2020). In the short run, new housing can affect existing units by causing a reallocation of people across submarkets.[1] Households who would have otherwise occupied cheaper units move into new units, reducing demand and prices for the housing they leave vacant. The process iterates when a second round of households moves into the units the first round left vacant, creating what is sometimes called a migration or vacancy chain. There is little empirical evidence on this short-run reallocation, although it is crucially important to policymakers seeking quick relief to declining housing affordability.

In this paper, I help fill this gap in the literature by using new data on individual address histories to directly study the migration chains started by new housing. My main exercises leverage this data to compute descriptive statistics that provide new insights into this long-hypothesized but empirically elusive mechanism. I then use a stylized simulation model to provide ballpark estimates of how these chains affect the low-income housing market. The paper also adds to scholarship on land-use regulation (reviewed by Gyourko and Molloy, 2015) by studying migration mechanisms and using a new empirical approach.[2]

My primary contribution is to document the sequences of moves sparked by new housing construction. I use address history data from Infutor Data Solutions to identify 686 large new market-rate multifamily buildings in 12 large central cities and track 52,000 of their current residents to their previous building of residence. I then find the tenants currently living in those buildings and track them to their pre-

☆ Thanks to Nathan Sotherland and Steve Yesiltepe for excellent research assistance. Thanks to Brian Asquith, George Galster, Atul Gupta, Ray Kluender, Stephen Nei, and Davin Reed for helpful comments. An earlier version of this paper previously circulated under the title "The Effect of New Luxury Housing on Regional Housing Affordability."

E-mail address: evanemast@gmail.com

[1] Grigsby (1963) was perhaps the first to formulate the theory of housing submarkets, and Rothenberg et al. (1991) further develop the idea into a model of a system of interconnected submarkets.

[2] This literature has generally used quasi-experimental (Glaeser, 2003; Glaeser and Ward 2009; Ihlanfeldt, 2007) or model-based (Nathanson, 2019; Anenberg and Kung, 2018) approaches.

https://doi.org/10.1016/j.jue.2021.103383
Received 2 February 2021; Received in revised form 16 July 2021
Available online 27 July 2021

vious residence (reweighting accordingly) and iterate for six rounds. When computing statistics about each round, I include only chains I can track to that round and, to prevent contamination from demographic differences across metros, only migrants from the same metropolitan area.

New building residents largely originate from nearby high-income areas. Sixty-seven percent arrive from within the same metro, and only 20% of those within-metro arrivals originate in below-median income tracts. However, the proportion from low-income areas rises steadily to 40% in round six, and similar patterns emerge for other neighborhood income or demographic thresholds. The chains also diffuse geographically—the sixth round has about the same suburban share as the metropolitan area as a whole. This exercise suggests that the reallocation resulting from new housing construction spans a diverse set of neighborhoods and is likely to loosen a wide spectrum of housing submarkets.[3]

My second contribution is a set of simulation estimates of how new high-end construction affects the low- and middle-income housing markets. The simulation is similar in spirit to the previous exercise but adds two key factors. First, instead of constructing the next round of a chain based on a household's origin unit, I use a calibrated estimate of where the household would have lived in a counterfactual world with no construction. For example, consider empty-nesters who moved from a suburb to a new building, but would have moved to some part of the city anyway. The next round of chain is the unit they would have moved to otherwise, not the suburban origin unit. Second, a chain can end if, for example, new construction induces household formation or migration from outside the metro.

Because these factors are unobservable and difficult to estimate, I run the simulation under a range of calibrations. The baseline, best suited to a marginal increase in supply, is the following. First, had the new buildings not been constructed, households would have lived in a neighborhood that is one decile higher in the metropolitan income distribution than their origin, but similar on other characteristics. Second, chains end due to unfilled vacancies resulting from matching frictions or landlord market power (proxied by neighborhood average vacancy rates), but there is no endogenous increase in out-of-metro migration or household formation in response to the new construction.

In the baseline, seventy percent of migration chains started by a new high-end unit reach a below-median income neighborhood before ending, implying that building 100 new units reduces demand for these areas by 70. Assuming that reducing demand and increasing supply have symmetric price effects, this can be interpreted as equivalent to adding 70 depreciated units in below-median income areas. In a more conservative specification that allows for some endogenous household formation and out-of-metro migration, more accurately capturing the effects of a larger supply increase, the percent of chains reaching below-median income areas falls to 45. For bottom-quintile income neighborhoods, these two calibrations return estimates of 40% and 17%. These estimates provide important context to the descriptive results—the observed migratory connectivity is strong enough that chains frequently reach low-income areas even under conservative assumptions. Moreover, the simulation implies that nearly all of the effect occurs within five years.

Finally, my third contribution is documenting migratory connectivity across neighborhoods more generally. Migration between neighborhoods with moderately different characteristics is quite common—for example, individuals originating in the fifth income decile frequently move to the fourth or seventh decile, but rarely the tenth or first. This helps reinforce the migration chain results, and it also suggests that other spatially concentrated shocks may diffuse across very different neighborhoods through a short series of common moves. This adds a new measure of connectivity within cities to a growing literature that has previously considered housing search (Piazzesi et al., 2020) and social networks (Bailey et al., 2020).

Together, these results suggest that new market-rate housing construction loosens the market for middle- and low-income housing, even in the short run, pointing to an important role for policies that increase construction. However, I do not estimate price effects, which are particularly unclear in neighborhoods where rents are already close to operating costs, leaving little room for reduced demand to lower them further. In addition, the endogenous response of out-of-metro migration and household formation may strengthen over time, moderating affordability effects in the long run.

2. Data

2.1. Infutor data

My primary data source is individual address histories from Infutor Data Solutions. Infutor constructs this information from various record sources—USPS change of address forms, county assessors, magazine subscriptions, et cetera. Addresses are reported at the building level (or even unit level in some cases) and are accompanied by an estimated arrival date. However, the data has poor coverage of people under age 25 and limited housing characteristics. I thus classify units based on their census tract characteristics, as measured in the 2013–2017 American Community Survey (ACS). Holding neighborhood characteristics fixed at their most recent values reduces the risk of misclassifying areas that gentrified during the sample period as low-income.

In the Appendix, I examine selection into the Infutor data. The data performs well, with the median tract containing 0.88 observations per individual over 25 reported in the Census and little correlation between coverage and demographic characteristics. Diamond et al. (2019) and Phillips (2020) provide further validation of the data.

2.2. New market-rate buildings

I identify new market-rate buildings using the Infutor data. I first collapse the individual-level data by street address and keep buildings with over 16 individuals (which, given Infutor's coverage rate, implies over 20 residents). I then identify new buildings as those where over 90% of current residents moved in since 2009 and keep only those that are within five miles of the central business district. In order to focus on the high-cost buildings that are most removed from the low-income market, as well as at the center of the policy debate, I keep only tracts that are above the metro median in either median household income or income per capita. Finally, I drop buildings—such as subsidized or student housing and homeless shelters—that meet the previous criteria but are not market-rate. The algorithm identifies both rental and owner-occupied buildings, as well as renovations of previously non-residential buildings. I include twelve large Core-Based Statistical Areas (CBSAs) with well-defined city centers: New York City, Chicago, Dallas, Houston, Washington, Philadelphia, Atlanta, Boston, San Francisco/Oakland, Denver, Seattle, and Minneapolis.

This yields 686 market-rate multifamily buildings (containing 52,432 individuals) constructed since 2009. Table 1 shows that all cities in the sample are well-represented and that buildings are in tracts with very high income and rent. Appendix Table A.1 shows some additional characteristics. The median building contains 60 Infutor individuals, and the buildings' tracts have relatively high vacancy rates: 10% on average and over 15% for many.

3. Migratory connections between neighborhoods

Before examining new housing directly, I study general migratory connections between neighborhoods. Because demographics vary

[3] These results broaden a small literature that computes similar statistics by sequentially interviewing a small sample of households (Kristof, 1965; Lansing et al., 1969) or examining administrative records in a single Scandinavian city (Turner, 2008; Turner and Wessel, 2019).

Table 1
Number of new buildings and residents across cities.

City	New buildings	Infutor individuals	Income decile	Rent decile	Percent from same CBSA	Percent from same City
Atlanta	44	3641	9.7	9.6	0.687	0.484
Boston	16	1238	9.8	9.9	0.700	0.375
Chicago	84	7068	10.0	9.9	0.728	0.578
Dallas	76	6670	8.9	9.8	0.687	0.487
Denver	49	3270	9.1	8.6	0.539	0.411
Houston	69	5906	9.9	9.4	0.711	0.584
Minneapolis	37	2206	9.6	9.5	0.714	0.483
New York City	89	7835	9.9	8.8	0.764	0.682
Philadelphia	12	694	9.9	9.9	0.642	0.423
Seattle	101	6334	9.6	9.5	0.598	0.472
San Francisco	38	1704	9.3	8.9	0.632	0.512
Washington	71	5866	9.2	9.7	0.658	0.505
Sample	686	52,432	9.6	9.5	0.672	0.500

Note: This table shows the number of new buildings in each city, the number of individuals currently living in those buildings in the Infutor data, mean characteristics of the tracts containing the buildings, and what percent of new building residents originated in the same city or CBSA. The buildings, which are detected using the algorithm described in Section 2, must contain over 16 individuals in the Infutor data, be built since 2009, and be within five miles of their CBSA's central business district and in a census tract with above median income for the CBSA. Tract characteristics are taken from the 2013–2017 ACS, and deciles are computed within each CBSA. Income is per capita, and rent is the median for a two-bedroom unit. Individuals whose immediately previous address is in the same CBSA (city) as the new building are considered from the same CBSA (city).

Table 2
Equivalent unit estimates.

		Simulated number of equivalent units in tract category				
Share marginal of new HHs, out-of-metro	Counterfactual step size	<P50 Inc.	<P50 Inc. & Rent Burdened	<P20 Inc.	<P20 Inc. & Rent Burdened	<P50 Inc. & <P20 White
Baseline calibration						
None	+1 inc. decile	0.702	0.453	0.396	0.258	0.488
Alternative decay rates						
25%	+1 inc. decile	0.453	0.191	0.167	0.088	0.231
Building-tract diff.	+1 inc. decile	0.513	0.239	0.204	0.110	0.279
Alternative step sizes						
None	+2 inc. decile	0.644	0.422	NA	NA	0.472
None	Data average	0.648	0.436	0.186	0.138	0.499
Alternative decay rates and step sizes						
Building-tract diff.	Data average	0.449	0.227	0.078	0.050	0.289

Note: This table shows the simulated effect of one new market-rate unit in a high-income area. The five rightmost columns show the expected number of equivalent units created in different neighborhood types, with each row representing a different calibration. The leftmost column shows the shares of new household formation and across-metro migration that are assumed to be marginal in each calibration. This is added to the submarket vacancy rate to construct the overall decay rate. "Building tract diff." is equal to the difference in the out-of-metro share of migrants to a new building and to its surrounding tract. The column second from the left provides the counterfactual step size calibration. "Data average" is equal to the average step size of out-migrants from a submarket (set to zero when the average is negative). An equivalent unit is created in a neighborhood type when a migration chain reaches it for the first time, as defined in Section 5.2. Tract characteristics are taken from the 2013–2017 ACS, and all quantiles are computed within CBSAs. Income is median household income, and rent burdened is defined as in the top quintile of rent burden for the CBSA.

greatly across cities, as shown in Appendix Table A.2, I focus on migration within the Chicago CBSA between 2010 and 2017.[4]

Fig. 1 plots distributions of destination tract characteristics conditional on origin tract characteristics. In Panel A, each box shows the median and interquartile range of destination median household income for migrants that originated in a given income decile. Whiskers show the 10th and 90th percentile. Individuals originating in top decile tracts very rarely move to a below-median neighborhood, and very few people from lower deciles migrate above the median. While this suggests that submarkets exist, they also appear to be permeable—individuals frequently move from the seventh decile to the ninth, the sixth to the fourth, et cetera. This implies that even top and bottom decile tracts can be connected through a short series of common moves. Panel B shows a similar pattern for two-bedroom rent.

Panel C shows results for percent white. The least white tracts are more separated from the remainder of the market than were either top or bottom income tracts. People outside of these tracts are extremely unlikely to move in. However, because individuals do move out of these areas, a connection to the broader housing market does exist.

[4] Appendix Fig. A.1 shows similar results for San Francisco.

These graphs show that even a metropolitan area that is highly segregated on most measures exhibits substantial migratory connections between tracts with very different characteristics. This suggests that migration chains from new housing could reach a wide variety of neighborhoods. More broadly, these results imply that geographically localized shocks are likely to have ripple effects on other parts of the region. For example, the effect of a new train stop will likely be highest in its immediate area, but it may also be large in places that have high cross-migration with that area, even if they are not nearby. Future research could use migratory connectivity to better understand the effects of localized shocks.

4. Migration chains from new housing

4.1. Method

In this section, I construct migration chains from new buildings. The exercise is conceptually simple. I start with the 52,000 individuals currently living in the 686 new market-rate buildings described in Table 1. I then use the Infutor data to identify their origin building. I then identify the people currently living in the first round's origin buildings and iterate for six rounds. This exercise is purely descriptive—it shows real-

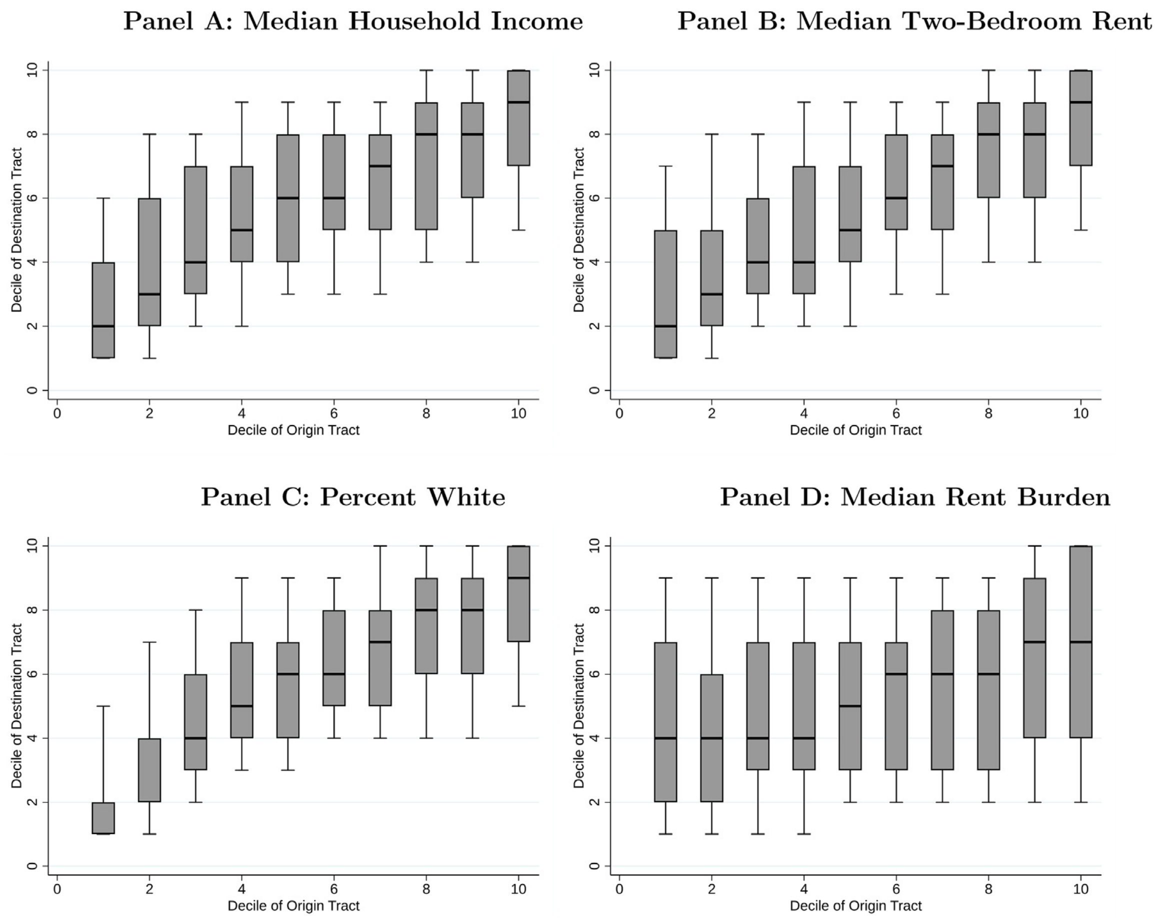

Fig. 1. Migration between census tracts in Chicago metropolitan area. *Note:* This figure shows the distribution of destination neighborhood characteristics conditional on origin neighborhood characteristics for migrants within the Chicago CBSA in 2010–2017. Within each panel, each box plot represents migrants who originated in a tract of a given decile of the characteristic in the heading. The box then shows the median and interquartile range of the same characteristic in the destination tracts of those migrants. The whiskers represent 10th and 90th percentiles. Characteristic deciles are calculated within the CBSA, and tracts that are over 20% undergraduate students are excluded.

location relative to before a building's construction, not relative to the counterfactual of no construction. It is also focused on migratory connectivity, not the rate at which chains end.

Despite the intuitive concept, there are several specification choices to note. First, after tracing a round of people to their origin residences, I keep only those that are within the same metro as the new building (about 70%, depending on the city and round). This prevents demographic differences across metros from polluting results.[5] Second, when identifying the people currently living in the previous rounds origin residences, I include all individuals in the origin buildings, rather than the specific origin units, then reweight so that they sum to one individual.[6] This both avoids inconsistencies in reported unit numbers that hinder matching and increases the probability that at least one person originated within the metro.

Finally, it is sometimes impossible to construct the next round of a chain. This can occur because I cannot track anyone in a building to an address within the same metro or because I cannot locate anyone currently living in a building vacated by a person in the previous round. In order to focus on connectivity rather than data imperfections or chain decay, I proportionally distribute the weight from the untracked building to other similar buildings in the round that can be tracked.[7] Sixty-seven percent of new building residents originated within the CBSA, and 74% of buildings in the next round have at least one person tracked within-metro. The figure gradually falls to 52% as single-family residences become more common.

4.2. Results

4.2.1. First round of chain

Fig. 2 illustrates the first round of the chain geographically by plotting new market-rate buildings in Chicago and the origin addresses of their current residents. New buildings are concentrated near the central business district and in wealthy neighborhoods to the north and northwest. However, residents of these buildings originated in a much broader area of the city. The highest-income areas send the most residents, but middle incomes are also well represented. Virtually no residents originated in zip codes with median income below $30,000. This

5 I also drop individuals from tracts that are over 20% undergraduate students.

6 To avoid major changes over time in building or neighborhood attractiveness, I restrict to people that moved into the building since 2009.

7 Similar buildings must have the same principal city/suburban status as the untracked building, which helps account for selection on tracking rates between multifamily and single-family buildings. Because it is computationally easier to compute the chain separately for each category in Fig. 3, the characteristics that tracked and untracked buildings are matched on change with each computation. In each case, I require similar buildings to have the same in/out of category status as the untracked building.

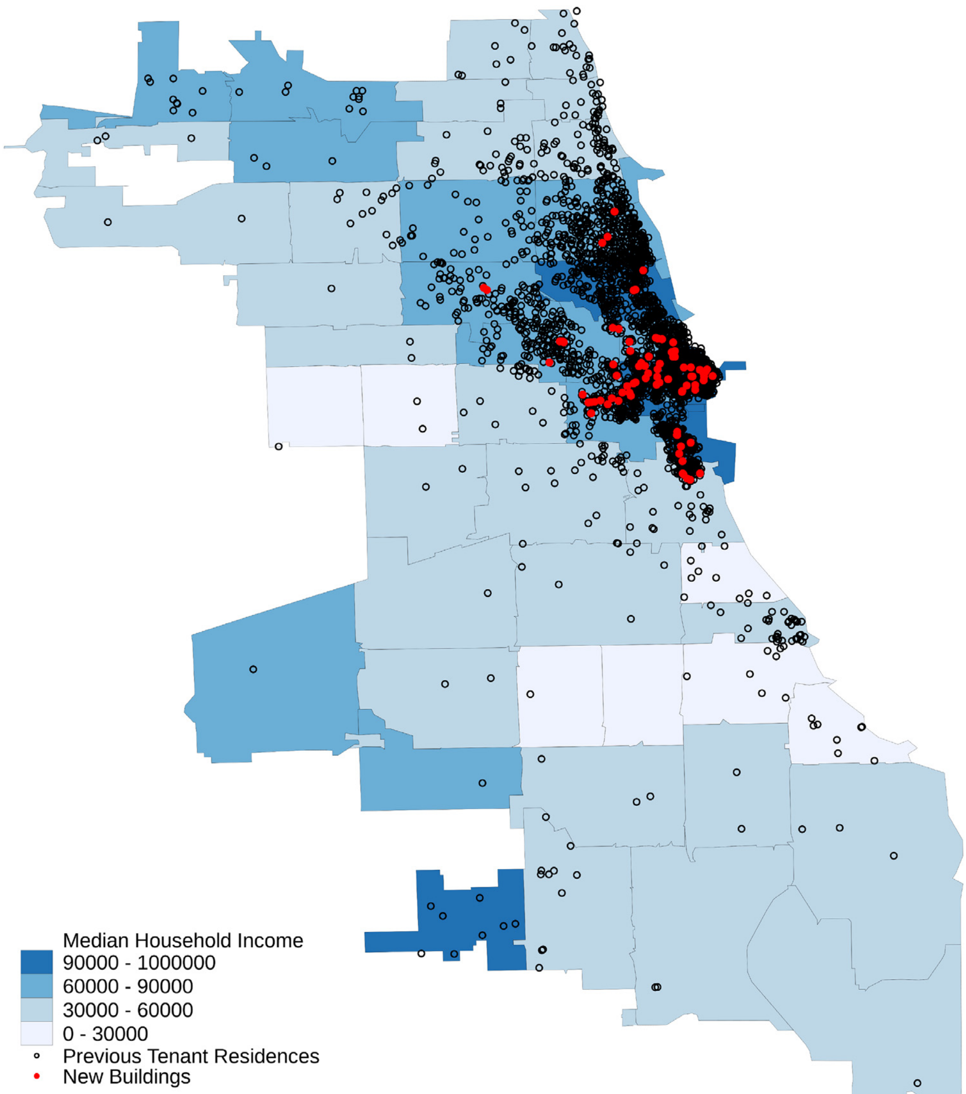

Fig. 2. Prior residence of tenants of new buildings in Chicago. *Note:* Solid red dots represent the location of market-rate multifamily buildings completed between 2010 and 2017. Hollow black dots represent the previous residences of the current tenants in those buildings. The base map polygons are zip codes in Chicago proper, colored according to median household income in the 2013–2017 ACS. Only residents whose prior residence was within the city proper are included. Small amounts of noise are added to each marker to avoid precisely identifying addresses. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

figure highlights that even when construction is highly concentrated, it will draw residents from large swathes of a city.[8]

Next, Table 1 shows that 67% of new building residents originate within the same CBSA, with 50% from the same city. This suggests that much of new buildings' affordability benefit will accrue locally and speaks directly to the policy debate, where many complain that new buildings primarily attract residents from other cities (Been et al., 2019). The percent of residents from the same CBSA is lowest in Denver (53.9%) and Seattle (59.8%), with the rest of the sample falling between 63% and 76%.

4.2.2. *Full chain*

Fig. 3 shows the percent of each round in five tract categories, defined according to the within-CBSA characteristic deciles shown in Appendix Table A.2 About 20% of new building residents originate in tracts with below-median median household income, rising to 40% by round six. The percent of individuals originating in a bottom-quintile income tract increases from 7% to 15% from the first to sixth round, and the percent from tracts that are below-median income and in the bottom quintile of percent white starts at six and eventually rises to 14.

These migration patterns are not consistent with strong segmentation in the housing market—even when tracing moves at the building level, it appears that a short series of moves connects new construction and low-income areas. To put these numbers in context, Appendix Fig. A.4 normalizes each line by the percent of CBSA population in each tract type. All categories are severely underrepresented in round 1—falling between roughly 15 and 40% of their prevalence in the CBSA. However, these numbers increase to between 70 and 80% by round six. Similarly, the percent of individuals originating within the principal city declines from 70% of round one to 30% in round six, close to the population average.

In addition to the main sample, I also repeat the exercise using the approximately 200 buildings that meet all of the sample criteria except the income restriction (Appendix Fig. A.5). Unsurprisingly, these sequences include a higher share of low-income units in every round. Appendix Fig. A.6 shows how results depend on the tract income decile of the new building. In the first round of the chain, the share of origin units in below-median income tracts is 55% for new buildings in the first income decile versus only 15% for buildings in the tenth decile. In the fifth round, the slope is more modest, as chains from both initially high- and low-income areas have converged towards the mean.

A potential issue with these exercises is that units are classified by tract characteristics, which may not match actual unit quality. To diagnose the extent of this problem, I first simply compare the tract and

[8] Appendix Fig. A.2 shows the Chicago CBSA, and Appendix Fig. A.3 repeats the exercise in San Francisco.

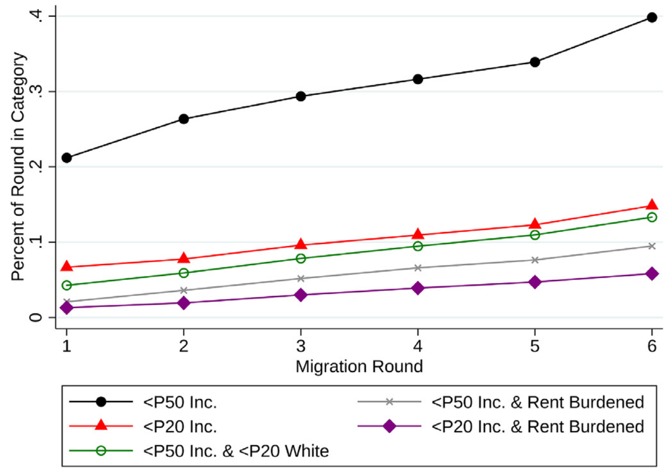

Fig. 3. Percent of individuals originating in tract categories by migration round. *Note:* This figure plots the percentage of the individuals in each round of the sequence of origin units that originated in a census tract with a given set of characteristics. Note that only migrants from the same metropolitan area as the new building are included in each round. Round 1 is the origin units of the individuals currently occupying the new unit; round 2, the origins of the individuals occupying round 1's origin buildings, and so on. Tract characteristics are taken from the 2013–2017 ACS, and all quantiles are computed within CBSAs. Income is median household income, and rent burdened is defined as in the top quintile of rent burden for the CBSA. Each subsequent round is constructed by observing the set of individuals currently living in the previous round's origin buildings, not their specific units, and the sequence is reweighted accordingly. The sequences begin with 52,000 individuals living in 686 new market-rate buildings.

block group characteristics of units in the sequence. The median difference between tract and block group median household income is only 0.5% in the first round and zero in the third round.[9] Second, I use the Infutor data to directly assess unit quality. I compare the likelihood that the average unit in, say, the fifth income decile is filled by a person moving from a lower income tract to the same probability for a fifth-decile unit that is included in the origin sequence. Panel A of Appendix Fig. A.8 shows that the units in the first round are somewhat less likely to be filled by a person from a lower income tract than the average unit in a similar tract. In the fifth income decile, the figure is about 30%, versus 36% in the full sample. However, as shown in Panel B, this gap falls to less than 1.5 percentage points by the third round. These two exercises suggest that while units in the sequence are somewhat positively selected, tract characteristics provide a reasonable description.

4.3. Discussion

These results show that the migration ripple effects of new housing reach a diverse set of neighborhoods, rather than only affecting high-income areas. While I do not measure price effects, past literature strongly suggests that increased vacancies will lower prices in their geographic area or submarket. In models of a thin housing market, a higher vacancy rate lowers both home prices and rents (Arnott, 1989; Wheaton, 1990). Recent empirical research also suggests that new apartment buildings lower nearby rents (Li, 2019; Asquith et al., 2019; Pennington, 2020) and increased sale listings lower nearby prices (Anenberg and Kung, 2014; Ater et al., 2021).

The results also add empirical depth to previous literature. Nathanson (2019) and Braid (1981) build models which generate reallocation across tiers of the market, and my statistics provide new evidence that these predictions are borne out in the data, not stopped by barriers between submarkets. More broadly, this migration underlies much of the large literature relating land-use regulation and housing prices.

5. Simulation model

5.1. Overview

In this section, I simulate migration chains that account for two important factors that moderate the effect on the low-income market and were absent from my previous exercises. First, individuals may have left their origin unit even if the new building was not constructed. For example, suppose a new building resident previously lived on Willow Drive, but would have moved from Willow Drive to Oak Lane if the new unit had not been constructed. Relative to the counterfactual of no construction, the new unit lowers demand for Oak Lane, not Willow Drive.

Second, chains may end with some probability in each round. A housing unit could be a second home or investment property, in which case the owner does not vacate their other unit. In addition, a new household could form to fill the new unit, or a household could move in from outside the metropolitan area. On the supply side, landlords with market power may not lower prices by enough to completely fill the vacancies. However, an important and empirically challenging complication is that chains only end if a household takes an action that they would not have in the no-construction counterfactual. For example, consider a person who moved to a new building in Philadelphia from a different metro. If they would not have moved to Philadelphia had the new building not been constructed, this ends the chain. If they would moved to Philadelphia regardless, the chain should instead proceed from the Philadelphia unit they would have occupied in the counterfactual.

Both of these forces are difficult to estimate, and prior literature provide limited guidance. Because of this, I use a range of calibrations to produce a set of ballpark estimates.

5.2. Model

The simulation produces a migration chain C, which is a sequence of housing units. C_1 is a new unit, and C_2 is the unit that the person living in the new unit would have occupied had the new unit never been constructed. I call this the person's counterfactual location. The sequence then iterates from C_2.

This recursive structure makes it easy to define the remainder of a chain given a starting point in a new building. To construct C_{i+1} given C_i, I first define O_i as the origin unit of the individual currently living in C_i. I then introduce the operator T, which maps an individual's origin unit to their counterfactual unit:

$$C_{i+1} = T(O_i).$$

In addition, I allow chains to end in each round with some probability $d(C_i)$.

In order to quantify the effect of the simulated chains, I say that a chain creates an "equivalent unit" in submarket k if it reaches that area before ending. The intuition is that inducing a household to move out of a submarket reduces demand for that submarket by one, which should have a similar effect on prices as building an additional (depreciated) unit in k. The expected number of equivalent units in k created by a new unit in h is just the probability that the chain reaches k. Formally,

$$EU_h(k) = P(C \cap k \neq \emptyset | C_1 \in h).$$

To estimate this quantity, I simulate a migration chain from each new building resident and then compute the probability empirically for a variety of submarkets k.[10] Note that the number of equivalent units in

[9] Appendix Fig. A.7 shows the full distribution of this difference.

[10] Because I do not observe household linkages in the Infutor data, I make the simplifying assumption that each individual occupies a distinct unit. This largely weights couples more heavily than singles, which may be appropriate given that the former generally occupy larger units.

different submarkets need not sum to one, because the same chain can hit multiple submarkets.

5.3. Calibration and estimation

Submarket definition. I begin by defining a set of submarkets. An address's submarket is determined by its tract's within-CBSA deciles of median household income and percent white, whether it is in the principal city, and whether it is in the top quintile of median rent burden. While I use the exact set of new buildings in the first round, the remainder of the simulation operates at the submarket level, which makes it easier to calibrate counterfactual locations systematically.

Baseline counterfactual location calibration. For migrants who originated within the same CBSA as the new building, the baseline calibration is that the counterfactual submarket is one income decile above their origin submarket and identical on other characteristics. This is motivated by Fig. 1, which suggests that the average migrant moves up by about one income decile. This may reflect the lifecycle trajectory of individuals climbing the housing ladder or negative net migration from low-income areas. Note that the simulation thus yields a set of submarkets for C_2, rather than exact addresses. I thus cannot construct O_2 by tracking the individuals in a specific set of addresses. Instead, I use the distribution of the origin submarkets of recent (2009–2017) arrivals to the C_2 submarket as the distribution of origin units in O_2 (and similarly in later rounds).

Out-of-metro migrants present a complication because the chain ends if their counterfactual location is also outside of the metro. Their treatment thus overlaps with assumptions on the chain decay rate. As described in detail below, the calibration of the decay rate will dictate that a certain percentage of out-of-metro arrivals were induced to move to the CBSA by the new building, ending the chain. For the arrivals that would have moved to the area anyway, I draw the counterfactual location from the distribution of counterfactual locations of within-metro migrants to the same new building (or submarket, depending on the round).[11]

Baseline decay rate calibration. My baseline calibration of d aims to capture the effect of a small increase to housing supply such as, say, one building with 100 units. I set the probability $d(C_1)$ of a chain ending in the first round equal to the vacancy rate in the block group containing C_1. In subsequent rounds, I use the average in C_i's submarket.

The intuition for this assumption is best illustrated by considering the new housing units, where the vacancy rate captures the units that are unfilled because landlords do not price them low enough or because they are not used as a primary residence (the ACS counts these as vacant), as well as matching frictions that cause some units to be vacant at a given time. Because a single building is unlikely to have a noticeable effect on prices that will induce more people to form households or move long distances, I assume that there will be no endogenous response on those margins. This logic fits well with a sequential search model in which people first decide whether to move CBSAs or start a new household, then choose a specific unit.

Alternative calibrations. I include estimates from two alternative counterfactual location assumptions. First, I simply set the step size to two income deciles instead of one. Second, I set the step size equal to the average step taken by out-migrants from each submarket. (In many high-income submarkets, the average observed step size is negative; I conservatively set the counterfactual step to zero in these cases.)

I also include calibrations with larger decay rates. These are better suited for large changes to housing supply that may increase household formation and migration across CBSAs or for longer-run effects, since endogenous responses likely strengthen over time. The 2018 Current Population Survey estimates that 11.5% of moves were to form a new household, and Table 1 shows that 32.8% of people in my sample of new units originated from outside the metropolitan area. First, I assume that these figures represent the average rates in each round of the chain and that new construction increases each force by 25%. This doubles the mean decay rate of 10% in the baseline. While this is not a formal upper bound, it assumes that new construction has a very large effect on these forces. Second, I replace this 25% figure with a more data-driven calibration—the difference between out-of-metro migration to new buildings and to their surrounding tracts. Specifically, for each new building b in tract t, I compute $\frac{out\_metro_b - out\_metro_t}{out\_metro_b}$, and I then take the average within each CBSA. This figure is slightly smaller—about 15%.

5.4. Results

Table 2 shows the simulated number of equivalent units under all calibrations. In the baseline, best suited to a small increase in supply, one hundred new market-rate units create 70.2 equivalent units in below-median income tracts. The estimates are also large for tracts that are even less similar to high-income areas, with 39.6 created in the bottom income quintile. Even for tracts that are below median income and in the bottom quintile of percent white, the figure is 48.8, though these areas are by definition a small percent of a CBSA. (Simultaneously presenting these estimates reiterates an important point—the equivalent metric should not be interpreted additively, as this double-counts ripple effects across submarkets.)

Results under higher decay rate calibrations, better suited for a large increase in supply, appear in rows 2 and 3 of Table 2. When across-metro migration and household formation increase by 25%, below-median income equivalent units fall to 45, a drop of 37% from the baseline. The number in the lowest income categories falls by more—59% in the bottom income quintile—because the effect of a higher decay rate increases exponentially in each round and these categories typically appear later in the chain. Estimates drop slightly less when basing the decay rate on the difference between building and tract out-of-metro arrival rates.

Results under two alternative counterfactual location assumptions are shown in rows 4 and 5. Increasing the step size to two income deciles decreases below-median income equivalent units from 70 to 64, and taking the average step size directly from the data yields a very similar effect. Finally, row 6 shows that using the building-tract difference decay rate and the step size from the data yields 45 and 8 below-median and bottom-quintile equivalent units, respectively.

These effects should be felt relatively quickly. In the baseline, nearly all below-median income equivalent units are created by round 15, and most in the bottom-quintile appear by round 20, as shown in Appendix Table A.3. To roughly estimate long a round takes, I use the average amount of time a rental unit is vacant and average time a home is on the market. These represent something of an upper bound because multiple rounds of a migration chain can happen simultaneously—suppose, for example, that one household that would have moved from Unit A to Unit B instead moves to the new building, leading another household searching at the same time to locate in Unit B instead of Unit C. Figures from Zillow and the Census Bureau imply that each round will take one to three months, suggesting that most effects will be felt within 2–5 years.[12]

5.5. Discussion

This range of estimates suggests that new construction will have a meaningful effect on the middle- and low-income housing market. Even under calibrations that imply that new construction creates very large

[11] I make the same assumption for individuals that are not tracked to their previous address in the Infutor data and individuals that originate in heavily college student (>20%) tracts.

[12] Zillow reports that the average house was on the market for one month before selling in 2018, though this number reached 140 days in 2010 (Zillow Group, 2020). The Census Bureau reports an average rental vacancy rate of 7% in 2018, with a peak near 11% in 2010. My estimate is also similar to Lansing et al. (1969), who found a mean and median lag time of 1.8 and 3.6 months, respectively.

changes in household formation and moves across metros, the simulations return significant equivalent unit figures. Unfortunately, because estimates of the elasticity of housing prices to increased quantities are limited, it is difficult to determine how the vacancies generated by the chains would affect prices. Appendix Table A.4 provides some back-of-the-envelope calculations of how submarket prices would change if enough high-end units were constructed to increase a metropolitan area's total housing stock by 10%.[13] Results for prices in below-median areas vary from −0.9% under the lowest elasticity of −0.1 (Li, 2019; Glaeser and Ward, 2009) to −6.3% under an elasticity of −0.7.

The estimates have implications for common housing affordability policies. Inclusionary zoning (IZ) requirements mandate that a certain percentage of new units must be income-restricted. While these policies can directly create affordable units, they also raise the cost of development, potentially slowing new construction and reducing indirect affordability benefits. I assess this trade-off with a simple back-of-the-envelope calculation. I assume a percentage point increase in the IZ requirement reduces the number of market-rate units constructed by x%, where x varies from 0 to 2. I then compute the number of total units, income-restricted units, and below-median income equivalent units that would be created under IZ requirements of 10% and 20%.[14] Results are shown in Appendix Table A.5. When x is below 1.27, increasing the IZ requirement from 10% to 20% increases the number of income-restricted units by more than it decreases equivalent units, generating a net increase in affordable units. Above this threshold, the opposite is true. While more research on the crowd-out effects of IZ is needed, this exercise suggests that there may be significant trade-offs to high affordability requirements. Of course, income-restricted units may offer benefits that private market vacancies do not.

6. Conclusion

The short-run effect of new market-rate housing on the market for middle- and low-income housing is crucial to the current policy debate, where government intervention and market-based strategies are often pitted against each other. My results suggest that new market-rate housing construction can improve housing affordability for middle- and low-income households, even in the short run. The effects are diffuse and appear to benefit diverse areas of a metropolitan area.

However, there are several shortcomings of market mechanisms. The most important may be in the lowest-cost and most rent-burdened submarkets. Census tracts that are in both the bottom quintile of median household income and the top quintile of rent burden have an average vacancy rate of 12.8%, compared to 8.1 in the rest of my sample. Given that rents are generally already low in such neighborhoods, this suggests that reducing demand through the migration chain mechanism is unlikely to lower costs further, perhaps because rents have reached the minimum cost of providing housing. In addition to potentially small price effects, there may also be important amenity effects reduced population in these areas, such as reduced retail options, school closures, or increased crime. However, the relationship between income and vacancy rates differs across cities—in New York City, vacancy rates in low-income and rent burdened tracts are 9.7 versus 8.8% in other tracts, while the figures are 20.8 and 8.4% in Chicago. Market mechanisms will likely be more effective at reducing prices in low-income areas that have low vacancy rates.

Declaration of Competing Interest

Author declare that he does has no conflict of interest.

[13] This increase is roughly the difference between construction in the San Francisco and Houston CBSAs over a ten-year period.

[14] While the targeted income range varies across cities, equivalent units in below-median income tracts are a reasonable comparison (Schuetz et al., 2009).

Appendix A. Data appendix

A1. Validation exercises

This appendix provides more information on the Infutor data's coverage and representativeness. Starting with raw population, the data closely tracks the Census over-25 population at the tract level, with a median of.88 observations per Census individual. The coverage rate is quite similar across demographic groups, as shown in Appendix Fig. A.9, which plots the ratio of the Infutor and Census populations against tract characteristics. The largest differences appear between tracts with different racial composition, with a coverage rate of about 80% in the least white tracts versus 95% in the whitest tracts.

Because my study primarily uses this data to track household migration, Infutor's coverage of moves is also important. The data misses a substantial number of moves—the annual individual migration rate in the Infutor data is 5.4%, compared to the 9.8% reported in the Census Bureau's 2018 Current Population Survey. This could occur because of difficulty linking individuals across moves, because of overlapping addresses, or because the Infutor data has poor coverage of highly mobile young adults. However, because my study uses each move separately, rather than following individuals through the lifecycle, it only requires that the moves that do appear in the Infutor data are randomly selected. To examine this, I next compute the average annual migration rate at the county level in the Infutor data and compare it to Census estimates (which are not available at the tract level). Appendix Fig. A.10 plots the ratio of the two estimates against county characteristics. Coverage appears to be relatively uncorrelated with county characteristics. There is a slight correlation with county income, where the ratio is about.38 in low-income counties and.45 in high-income counties. This correlation could lead to an underestimate of the number of individuals who move into a given neighborhood from a lower income neighborhood.

Another potential concern is that Infutor may not have the correct endpoints for moves. For example, if a household moves from A to B to C, the data could record the move as from A to C, overstating the connectivity between A and C. To examine this, I compute the difference in destination and origin county median household income for moves reported in Infutor and in the IRS Statistics of Income data. The distributions are extremely similar, as shown in Appendix Fig. A.11. In my simulation exercise, I also run robustness checks in which I remove moves across very different neighborhoods.

In addition to my exercises, further validation on this data set is now available in a number of academic papers, including Diamond et al. (2019) and Phillips (2020).

A2. Preparation of raw files

The initial Infutor data sets are an individual-level file with longitudinal address histories (the Consumer Reference Database, or CRD) and an individual-level file with the most recently observed address and a set of demographic variables (the Name and Address Resource Consumer, or NARC). Some demographics are directly observed and some are imputed, and the fields are populated with varying frequency. The data were provided by Infutor in spring 2018.

The CRD is the primary basis for the analysis, but the NARC is useful for two reasons. First, it helps to identify higher-quality observations. The raw CRD may include business addresses or something like a UPS store, while the NARC generally does not. I thus restrict to the subset of individuals that appear in both the NARC and the CRD. Second, the NARC has latitude and longitude information. I merge this information to the CRD to geocode addresses. Finally, I make two changes to the raw address sequences. I remove consecutive moves within the same building, which are often caused by missing unit numbers, and reorder address sequences to run from most to least recent.

A3. Appendix figures

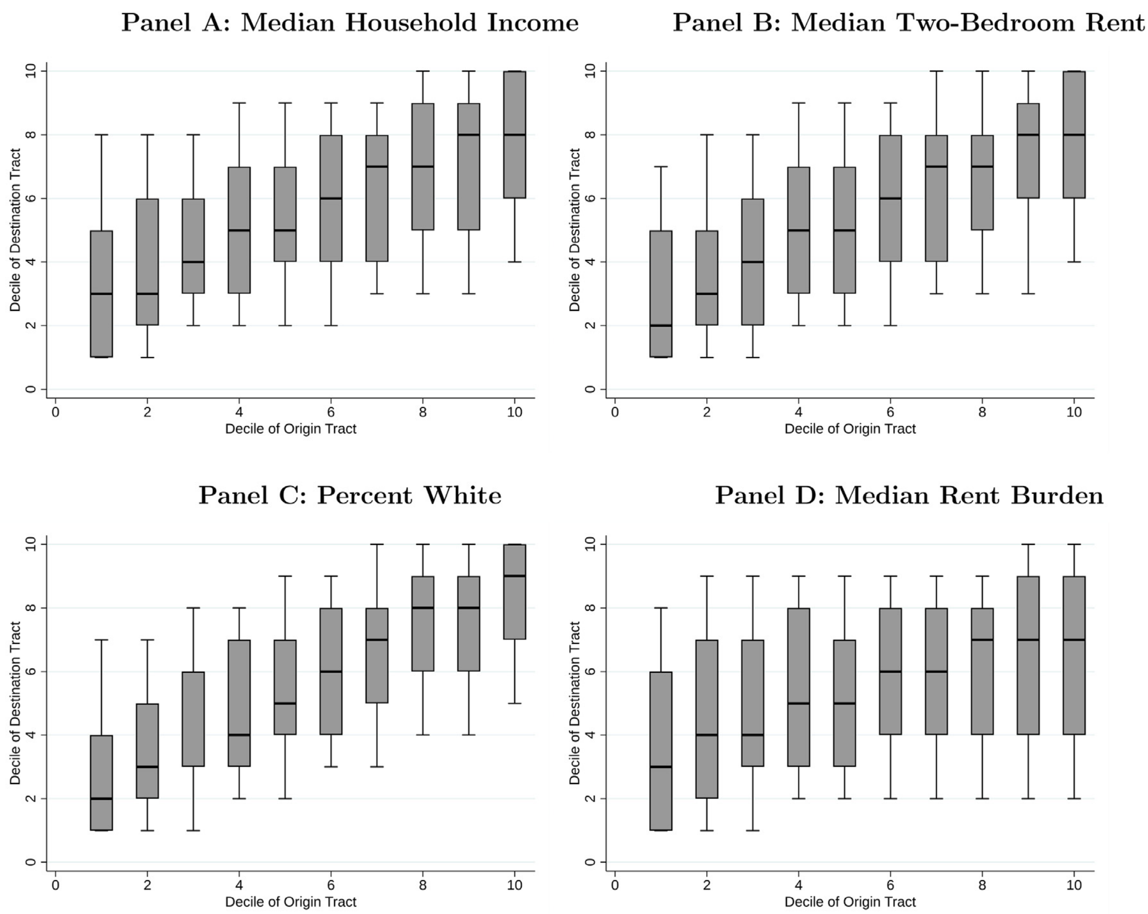

Fig. A.1. Migration between census tracts in San Francisco metropolitan area. *Note:* This figure shows the distribution of destination neighborhood characteristics conditional on origin neighborhood characteristics for migrants within the San Francisco CBSA in 2010–2017. Within each panel, each box plot represents migrants who originated in a tract of a given decile of the characteristic in the heading. The box then shows the median and interquartile range of the same characteristic in the destination tracts of those migrants. The whiskers represent 10th and 90th percentiles. Characteristic deciles are calculated within the CBSA, and tracts that are over 20% undergraduate students are excluded.

Table A.1
Building characteristics.

Percentile	Infutor individuals	Distance to CBD	Median household income decile	Income per capita decile	Median two-bedroom rent decile	Percent vacant (tract)	Percent vacant (block group)
Min	17	0.04	5	4	1	0.004	0.000
10	24	0.59	5	9	9	0.048	0.029
50	60	1.73	8	10	10	0.109	0.111
75	100	2.67	9	10	10	0.154	0.159
95	183	4.27	10	10	10	0.234	0.268
Max.	468	4.97	10	10	10	0.547	0.547
Mean	76.43	1.95	7.63	9.57	9.42	0.119	0.119
N	686	686	686	686	681	686	686

Note: This table shows characteristics of the new buildings. Distance to central business district (CBD) is given in miles. Median household income, income per capita, and median two-bedroom rent are determined using the building's census tract and the 2013–2017 ACS, and deciles of each are computed within CBSAs. Percent vacant is reported at both the tract and block group level and is also drawn from the ACS, which counts second homes as vacant.

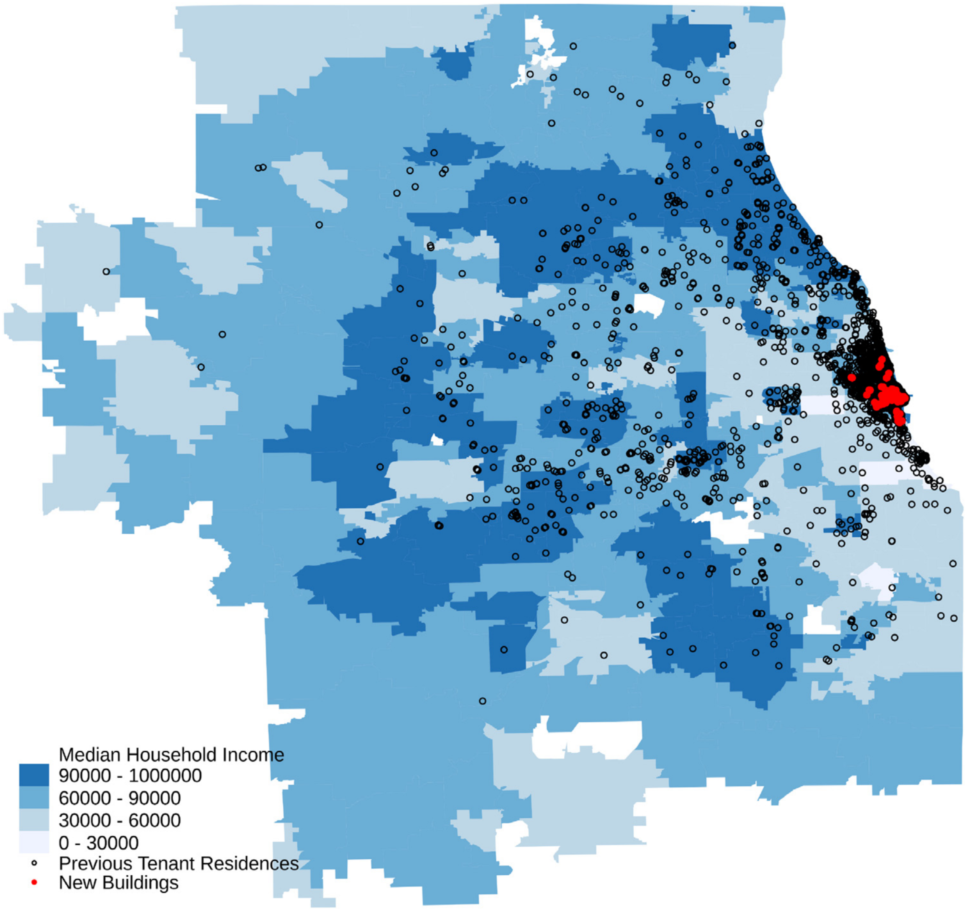

Fig. A.2. Chicago metro origins. *Note:* Solid red dots represent the location of market-rate apartment buildings completed between 2010 and 2017. Hollow black dots represent the previous residences of the current tenants in those buildings. The base map polygons are zip codes in the Chicago CBSA, colored according to median household income in the 2013–2017 ACS. Only residents whose prior residence was within the Chicago CBSA are included. Small amounts of noise are added to each marker to avoid precisely identifying addresses. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

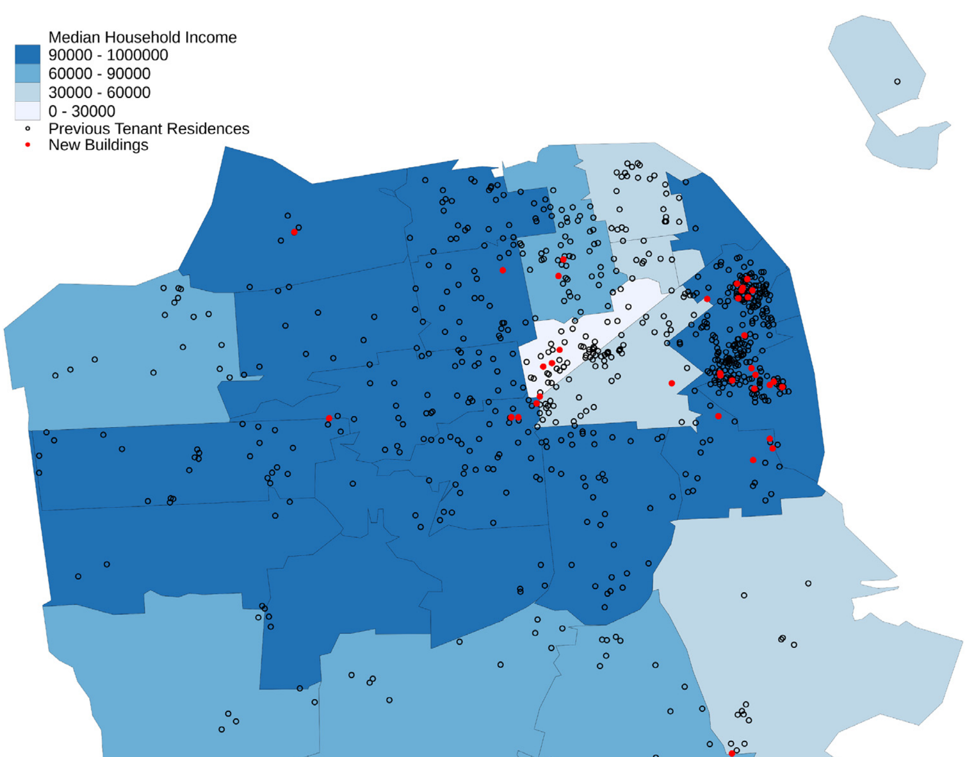

Fig. A.3. San Francisco city origins. *Note:* Solid red dots represent the location of market-rate apartment buildings completed between 2010 and 2017. Hollow black dots represent the previous residences of the current tenants in those buildings. The base map polygons are zip codes in San Francisco proper, colored according to median household income in the 2013–2017 ACS. Only residents whose prior residence was within the city proper are included. Small amounts of noise are added to each marker to avoid precisely identifying addresses. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Table A.2
CBSA characteristic deciles

MSA	Median income			Percent white			Median rent burden		
	P20	P50	P90	P20	P50	P90	P20	P50	P90
Atlanta	41,375	61,496	108,295	14.9%	51.0%	85.7%	24.7%	29.5%	40.1%
Boston	57,613	83,865	135,858	49.9%	78.9%	94.8%	25.2%	29.7%	39.9%
Chicago	40,024	62,601	113,019	9.5%	56.8%	87.6%	24.5%	30.1%	44.6%
Dallas	40,733	61,300	118,241	19.1%	48.8%	82.0%	24.3%	28.7%	38.1%
Washington	64,810	95,690	160,833	18.5%	50.4%	83.0%	24.2%	28.8%	40.2%
Denver	49,918	72,031	117,917	43.6%	73.3%	89.2%	24.9%	29.6%	38.5%
Houston	36,932	57,136	112,357	7.7%	34.5%	73.8%	24.2%	29.1%	39.1%
Minnneapolis	52,471	72,357	111,406	62.0%	81.4%	94.5%	24.0%	28.6%	37.7%
New York	45,677	72,657	129,479	8.1%	50.1%	87.1%	26.7%	32.3%	46.8%
Philadelphia	41,897	68,152	116,053	28.1%	74.0%	92.3%	25.0%	30.7%	46.0%
Seattle	55,530	79,040	117,500	50.9%	68.8%	86.0%	24.7%	29.2%	37.3%
San Francisco	62,731	96,210	157,045	17.5%	40.6%	74.0%	24.2%	29.2%	39.2%

Note: This tale shows CBSA deciles of the characteristics used to define submarkets. Characteristics are drawn from the 2013–2017 ACS.

Table A.3
Equivalent unit totals by round.

Round	Number of equivalent units in:				
	<P50 Inc.	<P50 Inc. & Rent Burdened	<P20 Inc.	<P20 Inc. & Rent Burdened	<P50 Inc. & <P20 White
5	0.539	0.186	0.164	0.077	0.238
10	0.659	0.313	0.262	0.141	0.361
15	0.690	0.379	0.318	0.181	0.422
20	0.699	0.414	0.350	0.208	0.453
25	0.701	0.432	0.369	0.225	0.470
30	0.702	0.442	0.380	0.236	0.478
35	0.702	0.447	0.386	0.244	0.482
100	0.702	0.453	0.396	0.258	0.488

Note: This figure shows the expected number of equivalent units created by a new market-rate unit, cumulative across rounds of the migration chain. An equivalent unit is created in, for example, a below-median income submarket when a migration chain reaches such an area for the first time, thus reducing demand for that submarket by one. Tract characteristics are taken from the 2013–2017 ACS, and all quantiles are computed within CBSAs. Income is median household income, and rent burdened is defined as in the top quintile of rent burden for the CBSA.

Table A.4
Submarket price effects of 10% increase in housing stock.

Submarket	Elasticity			
	0.1	0.3	0.5	0.7
Below-median income	0.9%	2.7%	4.5%	6.3%
Bottom-quintile income	0.8%	2.5%	4.2%	5.8%
Above-median income	2.0%	6.0%	10.0%	14.0%

Note: This table presents back-of-the-envelope calculations of the change in submarket prices caused by increasing the total housing stock by 10% through high-end construction. Estimates are produced by computing the number of equivalent units that would be created in the submarket of interest (assuming that the housing stock is proportionally distributed across submarkets) and multiplying by the elasticity given in the column header. Because this is a large change in supply, I use the conservative equivalent unit estimates from the calibration that assumes that 25% of across-metro migration and household formation are marginal to the new construction. The elasticity parameter is the percent change in submarket prices resulting from a 1% increase in submarket housing stock.

Table A.5
Effect of inclusionary zoning on income-restricted units and equivalent units.

	IZ effect on market-rate units				
	0	1	1.27	2	
10% IZ requirement	111	100	97	88.9	*Total units*
	100	90	87.3	80	*Market-rate*
	11.1	10	9.7	8.89	*Income-restricted*
	70.2	63.2	61.3	56.2	*Below-median EU*
	81.3	73.2	71.0	65.0	*Total affordable*
20% IZ requirement	125	100	93.3	75	*Total units*
	100	80	74.6	60	*Market-rate*
	25	20	18.7	15	*Income-restricted*
	70.2	56.2	52.4	42.1	*Below-median EU*
	95.2	76.2	71.0	57.1	*Total affordable*

Note: This table presents back-of-the-envelope calculations on the number of market-rate units, income-restricted units, and below-median income equivalent units that would be created under inclusionary zoning (IZ) requirements of 10% and 20%. For simplicity, I assume that 100 market-rate units would be constructed with no IZ requirement. I use the baseline estimate of equivalent unit creation and a range of calibrations for the effect of IZ requirements on new market-rate construction. For the latter, each column represents the percent reduction in market-rate units caused by a one percentage point increase in the IZ requirement. For example, column 3 assumes that a 1 percent increase in the requirement will reduce the number of market rate units by 1.4 percent.

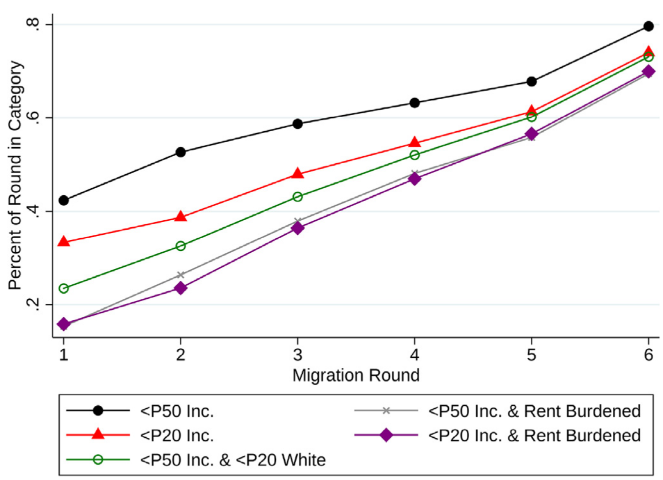

Fig. A.4. Normalized composition of sequence of origin units. *Note:* This figure repeats Fig. 3, normalizing each line by the percent of the CBSA population that lives in a given tract type.

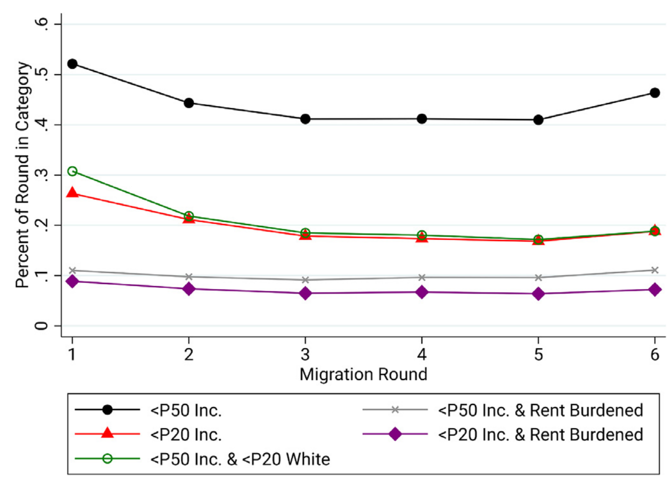

Fig. A.5. Percent of individuals originating in tract categories by migration round (new buildings in low-income areas). *Note:* This figure repeats Fig. 3 using new buildings in low-income tracts as the starting point, rather than the buildings in high-income areas that are included in the main sample. The buildings included meet all of the main sample requirements except the income restriction.

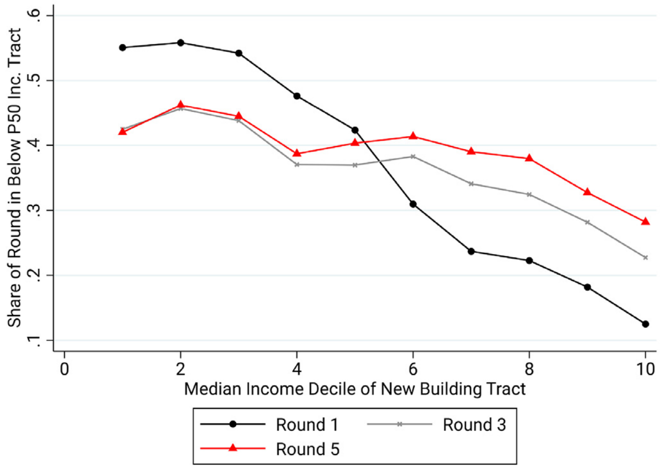

Fig. A.6. Percent of individuals originating in below-median income tracts versus tract income of new building. *Note:* This figure plots the share of origin sequence units in below-median income tracts against the characteristics of the new building's tract. Each line represents a different round of the sequence. New buildings are separated based on the income decile of their tract. The underlying estimation is identical to the main aggregate results shown in Fig. 3.

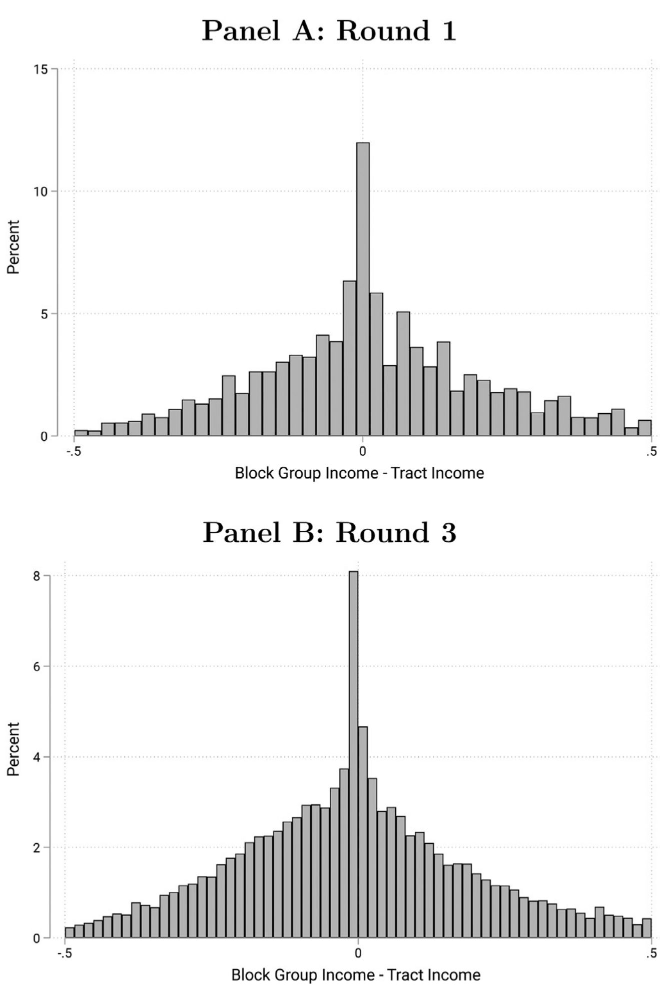

Fig. A.7. Distribution of block group income minus tract income in origin sequence. *Note:* These histograms show the distribution of the difference between tract median household income and block group median household income for units in the sequence shown in Fig. 3. The top panel plots this for units in the first round of the chain, and the bottom shows the third round. Characteristics are drawn from the 2013–2017 ACS.

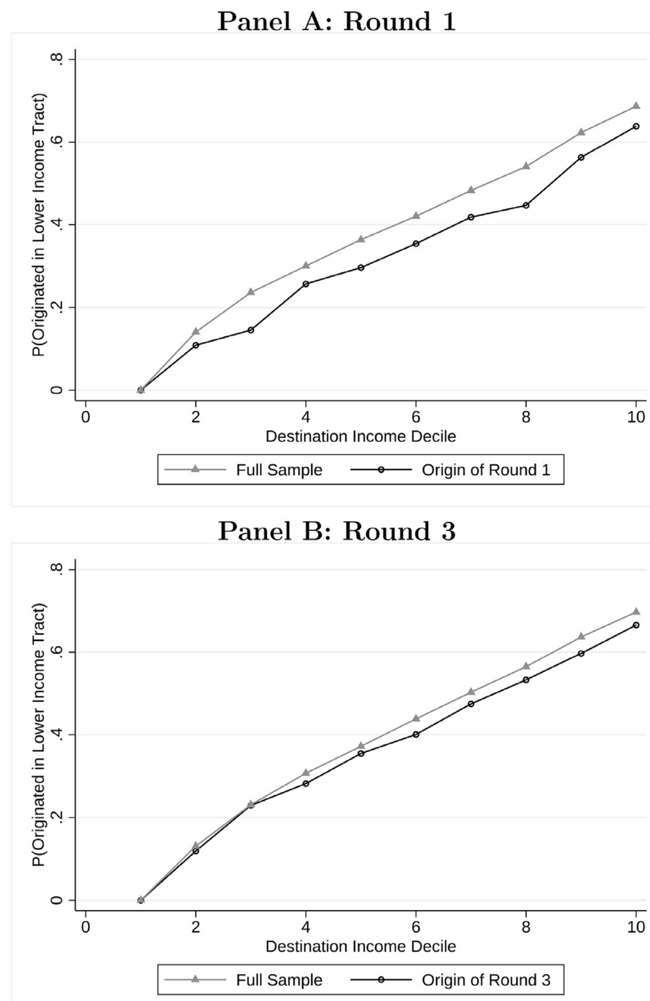

Fig. A.8. Percent of units filled by individual from lower income decile for full sample versus sequence of origin units. *Note:* Each line shows the probability that a unit in a given tract income decile was filled by a person who originated in a lower income tract. The full sample line includes all units in the income decile, while the other line includes only units that were in round 1 (Panel A) or round 3 (Panel B) of the sequence of origin units.

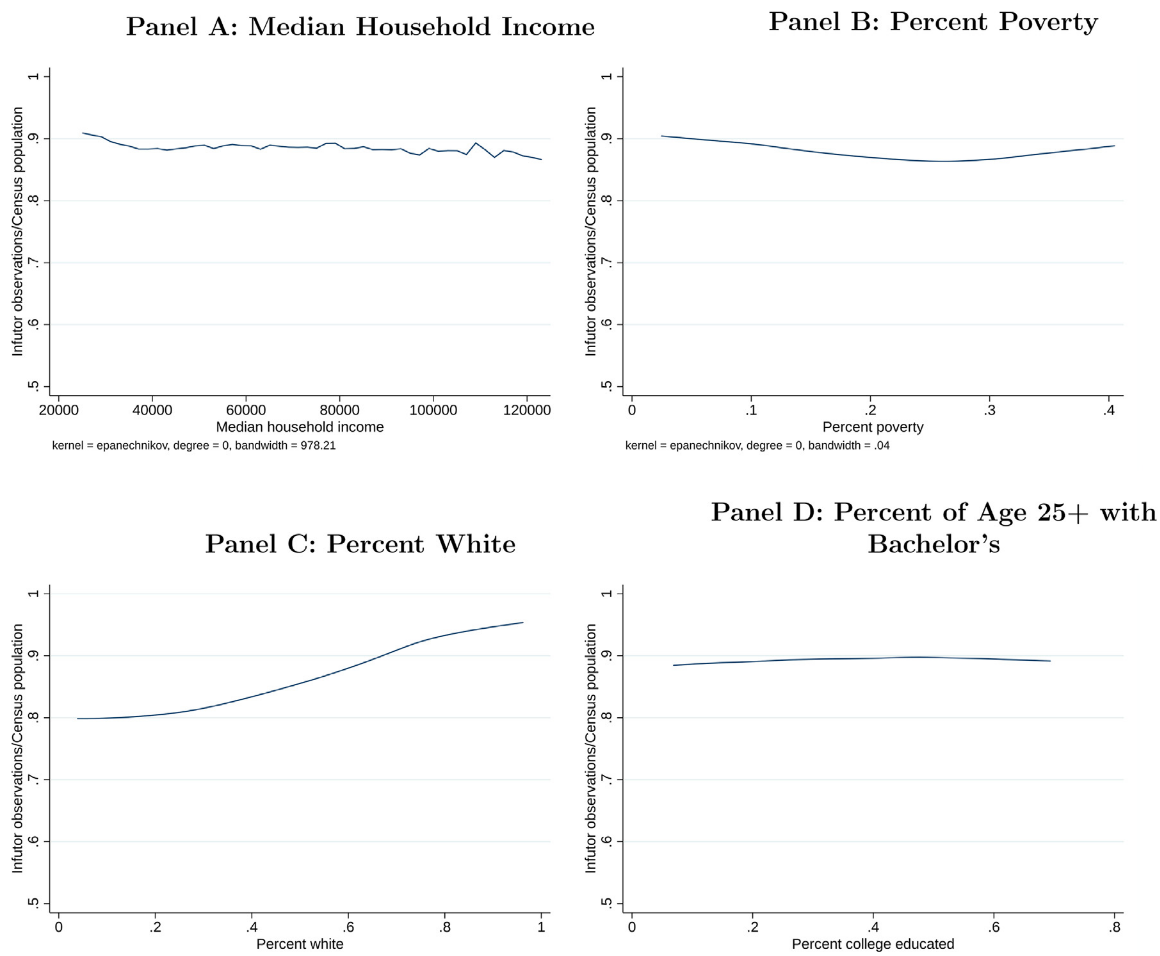

Fig. A.9. Infutor vs. census population (census tract level). *Note:* Each panel plots a local polynomial regression of Infutor coverage (measured as the ratio of Infutor observations to Census over-25 population) in a census tract versus the tract characteristic in the heading. Tract characteristics are drawn from the 2013–2017 ACS.

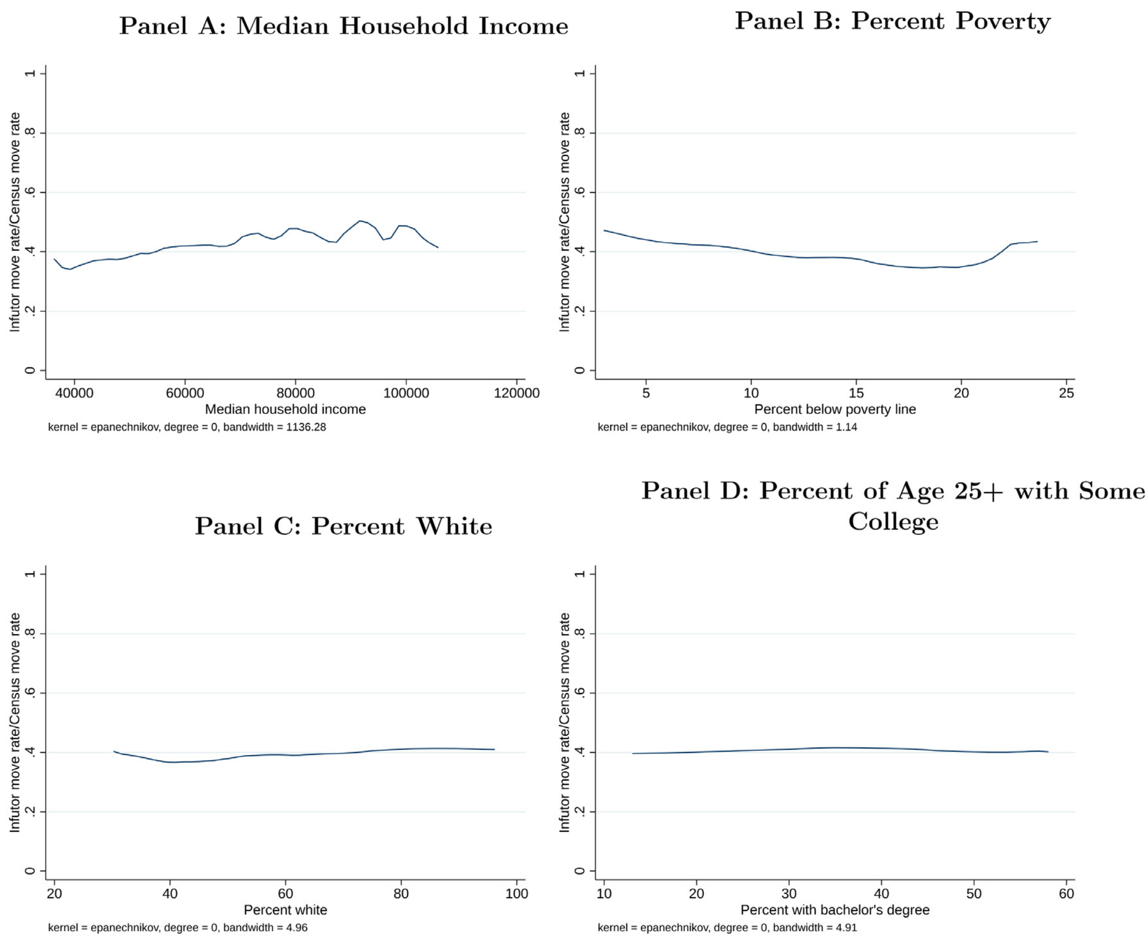

Fig. A.10. Infutor vs. census migration rates (county level). *Note:* Each panel plots a local polynomial regression of the ratio of Infutor to Census annual move rates (measured at the county level) against county characteristics. County characteristics and move rates are drawn from the 2013–2017 ACS.

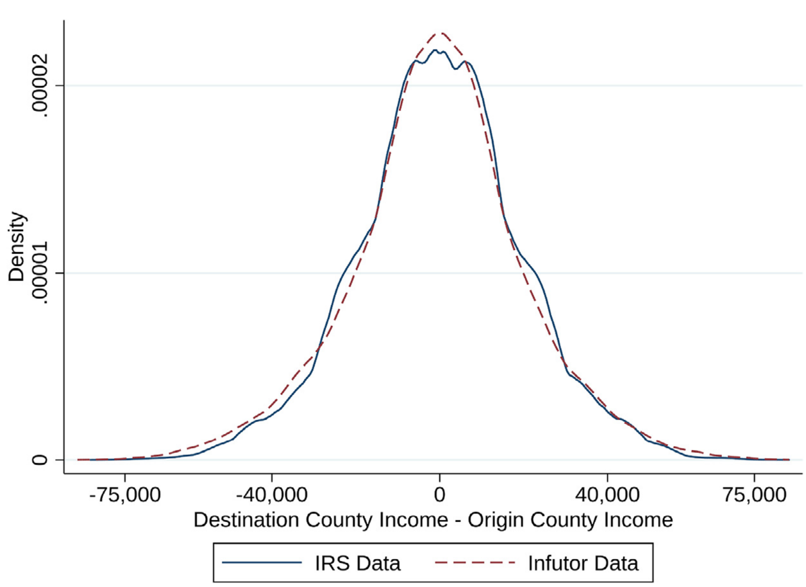

Fig. A.11. Difference between destination and origin county income in Infutor versus irs data. *Note:* This figure plots the distributions of destination county income - origin county income for moves in the Infutor data and 2018 IRS Statistics of Income data. County median household income is taken from the 2013–2017 ACS.

References

Anenberg, E., Kung, E., 2014. Estimates of the size and source of price declines due to nearby foreclosures. Am. Econ. Rev. 104 (8), 2527–2551.

Anenberg, E., Kung, E., 2018. Can more housing supply solve the affordability crisis? Evidence from a neighborhood choice model. Reg. Sci. Urban Econ. 80, 103363.

Arnott, R., 1989. Housing vacancies, thin markets, and idiosyncratic tastes. J. Real Estate Finance Econ. 2 (1), 5–30.

Asquith, B., Mast, E., Reed, D., 2019. Supply shock versus demand shock: the local effects of new housing in low-income areas. In: Upjohn Institute Working Paper, pp. 19–316.

Ater, I., Elster, Y., Hoffman, E., 2021. Real-estate investors, house prices, and rents: evidence from capital-gains tax changes. Mimeo.

Bailey, M., Farrell, P., Kuchler, T., Stroebel, J., 2020. Social connectedness in urban areas. J. Urban Econ. 118, 103264.

Been, V., Ellen, I.G., O'Regan, K., 2019. Supply skepticism: housing supply and affordability. Hous. Policy Debate 29 (1), 25–40.

Braid, R.M., 1981. The short-run comparative statics of a rental housing market. J. Urban Econ. 10 (3), 286–310.

Diamond, R., McQuade, T., Qian, F., 2019. The effects of rent control expansion on tenants, landlords, and inequality: evidence from San Francisco. Am. Econ. Rev. 109 (9), 3365–3394.

Glaeser, E.L., Ward, B.A., 2009. The causes and consequences of land use regulation: evidence from Greater Boston. J. Urban Econ. 65 (3), 265–278.

Glaeser, E.L., 2003. The impact of building restrictions on housing affordability (joint with J. Gyourko). Federal Reserve Bank of New York Economic Policy Review 9 (2), 21–39.

Grigsby, W.G., 1963. Housing markets and public policy. University of Pennsylvania Press.

Gyourko, J., Molloy, R., 2015. Regulation and housing supply. In: Handbook of Regional and Urban Economics, 5. Elsevier, pp. 1289–1337.

Ihlanfeldt, K.R., 2007. The effect of land use regulation on housing and land prices. J. Urban Econ. 61 (3), 420–435.

Kristof, F.S., 1965. Housing policy goals and the turnover of housing. J. Am. Inst. Plan. 31 (3), 232–245.

Lansing, J.B., Clifton, C.W., Morgan, J.N., 1969. New Homes and Poor People: A Study of Chains of Moves. Institute for Social Research, University of Michigan, Ann Arbor, MI.

Li, X., 2019. Do new housing units in your backyard raise your rents?, Mimeo.

Liu, L., McManus, D. A., Yannopoulos, E., 2020. Geographic and temporal variation in housing filtering rates. Mimeo.

Nathanson, C. G., 2019. Trickle-down housing economics. Mimeo.

Pennington, K., 2020. Does building new housing cause displacement?: the supply and demand effects of construction in San Francisco. Mimeo.

Phillips, D.C., 2020. Measuring housing stability with consumer reference data. Demography 57 (4), 1323–1344.

Piazzesi, M., Schneider, M., Stroebel, J., 2020. Segmented housing search. Am. Econ. Rev. 110 (3), 720–759.

Rosenthal, S.S., 2014. Are private markets and filtering a viable source of low-income housing? Estimates from a "repeat income" model. Am. Econ. Rev. 104 (2), 687–706.

Rothenberg, J., Galster, G.C., Butler, R.V., Pitkin, J.R., 1991. The Maze of Urban Housing Markets: Theory, Evidence, and Policy. University of Chicago Press.

Schuetz, J., Meltzer, R., Been, V., 2009. 31 flavors of inclusionary zoning: comparing policies from San Francisco, Washington, DC, and Suburban Boston. J. Am. Plan. Assoc. 75 (4), 441–456.

Sweeney, J.L., 1974. A commodity hierarchy model of the rental housing market. J. Urban Econ. 1 (3), 288–323.

Turner, L.M., 2008. Who gets what and why? Vacancy chains in Stockholm's housing market. Eur. J. Hous. Policy 8 (1), 1–19.

Turner, L.M., Wessel, T., 2019. Housing market filtering in the oslo region: pro-market housing policies in a nordic welfare-state context. Int. J. Hous. Policy 19 (4), 1–26.

Weicher, J.C., Eggers, F.J., Moumen, F., 2010. The Long-Term Dynamics of Affordable Rental Housing. Hudson Institute, Washington, DC.

Wheaton, W.C., 1990. Vacancy, search, and prices in a housing market matching model. J. Polit. Econ. 98 (6), 1270–1292.

Zillow Group, 2020. What is the average time to sell a house? https://www.zillow.com/sellers-guide/average-time-to-sell-a-house/.