Supply Shock Versus Demand Shock: The Local Effects of New Housing in Low-Income Areas

Brian J. Asquith
W.E. Upjohn Institute for Employment Research

Evan Mast
W.E. Upjohn Institute for Employment Research

Davin Reed
Federal Reserve Bank of Philadelphia

February 2020

FEDERAL RESERVE BANK OF PHILADELPHIA

Ten Independence Mall, Philadelphia, PA 19106-1574 • (215) 574-6458 • www.philadelphiafed.org/community-development/

Supply Shock Versus Demand Shock: The Local Effects of New Housing in Low-Income Areas*

Brian J. Asquith† Evan Mast‡ Davin Reed§

Abstract

We study the local effects of new market-rate housing in low-income areas using microdata on large apartment buildings, rents, and migration. New buildings decrease nearby rents by 5 to 7 percent relative to locations slightly farther away or developed later, and they increase in-migration from low-income areas. Results are driven by a large supply effect—we show that new buildings absorb many high-income households—that overwhelms any offsetting endogenous amenity effect. The latter may be small because most new buildings go into already-changing areas. Contrary to common concerns, new buildings slow local rent increases rather than initiate or accelerate them.

JEL Codes: R21, R23, R31
Keywords: Housing supply, housing affordability, gentrification, amenities

*We thank Jan Brueckner, Joshua Clark, Lei Ding, Atul Gupta, Andrew Hanson, Ray Kluender, Xiaodi Li, Jeffrey Lin, Otis Reid, and Jenny Schuetz for helpful comments and suggestions. We also thank Zillow™ for sharing data on rental listings and Real Capital Analytics for sharing data on new rental building locations. Shane Reed, Nathan Sotherland, and Steve Yesiltepe provided excellent research assistance. This Philadelphia Fed working paper represents preliminary research that is being circulated for discussion purposes. The views expressed in this paper are solely those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Philadelphia or the Federal Reserve System. Any errors or omissions are the responsibility of the authors. Philadelphia Fed working papers are free to download at: https://philadelphiafed.org/research-and-data/publications/working-papers.

†Asquith: W.E. Upjohn Institute for Employment Research. asquith@upjohn.org

‡Mast: Corresponding author. W.E. Upjohn Institute for Employment Research. mast@upjohn.org

§Reed: Federal Reserve Bank of Philadelphia, Community Development and Regional Outreach Department. davin.reed@phil.frb.org

Introduction

Housing costs have risen rapidly relative to incomes over the past 60 years in the United States, particularly in large and economically successful cities (Gyourko et al. 2013; Albouy, Ehrlich, and Liu 2016). This trend has increased rent burdens for low-income households, reduced regional economic convergence, and slowed national economic growth (Ganong and Shoag 2017; Hsieh and Moretti 2019). One potential solution is to increase the supply of housing by allowing more market-rate construction, which both theory and recent empirical results suggest will reduce regional housing costs (Rosenthal 2014; Mast 2019).

However, new housing developments could counterintuitively increase costs in their immediate area, raising questions about the incidence of such policies. This could occur because of amenity or signaling effects—if new units attract high-income households and new amenities that make the area more appealing, it could raise demand by enough to offset the increased supply. Recent literature on the importance and endogeneity of amenities suggests that this mechanism is plausible.[1] Moreover, worries about local rent increases lead many residents and policymakers to strongly oppose new construction, especially in low-income or gentrifying areas where high-end apartments and their occupants represent a large change from the status quo (Hankinson 2018; Been, Ellen, and O'Regan 2019). Although this idea has played a significant role in the housing debate, the very local effects of new market-rate housing have been understudied, likely due to data limitations and the endogenous location of new construction.[2]

This paper provides new empirical evidence on how the construction of large market-rate rental apartment buildings in low-income, central city neighborhoods affects nearby rents and migration. We assemble granular, address-level microdata on new construction and outcomes for 11 major cities, which allow us to address endogeneity concerns by leveraging extremely local variation in the timing and location of new construction.[3] We find that new buildings lower nearby rents by 5 to 7 percent relative to trend and increase in-migration from low-income areas. If there is an endogenous amenity effect, it appears to be overwhelmed by the standard supply effect. Our results suggest that, unlike in other contexts where diffuse policy benefits must be weighed against concentrated costs, there is not a trade-off between

[1] Guerrieri, Hartley, and Hurst (2013); Diamond (2016); and Su (2019) demonstrate the importance of endogenous amenities and residential sorting for housing costs and neighborhood composition. Baum-Snow and Marion (2009), Diamond and McQuade (2019), and Davis et al. (2019) show that building Low-Income Housing Tax Credit housing can increase home prices in some cases.

[2] For example, opposition from tenant groups concerned about rent increases near new development helped defeat California Senate Bill 50, which aimed to broadly increase residential zoning (Brey 2019; Dillon 2019).

[3] Our sample includes Atlanta, Austin, Chicago, Denver, Los Angeles, New York City, Philadelphia, Portland, San Francisco, Seattle, and Washington, D.C.

the regional and local effects of new housing construction. It benefits both.

Causal identification in this setting is challenging because developers select the locations of new buildings based in part on unobserved local characteristics and trends. In addition, the size and shape of a new building's amenity or reputation effects is unknown, making it difficult to know where they may shrink or reverse the negative effect of added supply. We attempt to overcome these challenges by leveraging our unique data to construct three related empirical strategies. The first is a difference-in-differences specification that compares the area very close to a new building to the area slightly farther away (our "near-far" specification).[4] The idea is that frictions in the land assembly and development approval processes lead to random variation in building placement and timing at the hyper-local level, making the outer area a good control for the treated inner area. This specification is well suited to detect one way that new buildings could raise rents—through amenity effects that fade out quickly with distance, such as increased retail options or the replacement of a vacant lot.

The second exercise is a difference-in-differences that compares listings near buildings completed in 2015 and 2016 to listings near buildings completed in 2019, after the conclusion of our sample (our "near-near" specification). The underlying logic is that developers choose sites in both groups for similar reasons, but one building is completed before the other for idiosyncratic reasons, such as the timing of when sites are available for purchase. Because the treatment and control areas are not necessarily in the same neighborhood, this specification can detect price changes driven by broader effects that may cross the near-far boundary of the first exercise. A variant of this specification also allows us to examine how congestion effects near the building could influence our results. Finally, we combine both sources of variation into a triple-difference specification that effectively compares the near-far difference around 2015–2016 buildings to the near-far difference around 2019 buildings.

We first study the effect on rents using listing-level data from Zillow™ that span 2013 to 2018. All three empirical approaches show that new construction in low-income neighborhoods (census tracts)[5] reduces nearby rents by 5 to 7 percent.[6] Event study plots support the parallel trends assumption and suggest that the negative effect begins in the same year as building completion and persists for at least three years after. The most likely source of

[4]Our primary specification uses 250 meters as the treatment group and 250–600 meters as the control group. The former is one or two city blocks, and 600 meters is an 8- to 10-minute walk. Schwartz et al. (2006) provide evidence that the positive effect of replacing blighted lots fades quickly within this range.

[5]We define low-income neighborhoods as census tracts with median household incomes below the metropolitan area median, which is a common definition of "gentrifiable" used in the gentrification literature.

[6]We consider the sample of new buildings in all neighborhoods in an extension and find noisy estimates that are statistically indistinguishable from zero. This may occur because the demand elasticity for established high-income areas is much larger than for gentrifying low-income areas, leading a supply shock to have a smaller effect on prices.

remaining bias—that developers still have some ability to target the best locations and times even at a very local scale—would shift our estimates toward zero.

While our strategies are well-suited to detect rent increases driven by amenity effects that center around the new building, other shapes are possible. For instance, if buildings generate congestion effects that fade out very quickly but relatively broad positive amenity effects, rents could increase in a doughnut shape around the building. We test this story directly and find no supporting evidence. In an extension, we use the empirical derivative method of Diamond and McQuade (2019) to estimate a continuous rent effect of new buildings and find that it shrinks monotonically with distance. While we cannot test every possible spatial pattern in which new buildings could increase prices, our exercises rule out the most likely alternative stories.

In our second set of results, we study the effect on in-migration using individual address histories from Infutor Data Solutions. In-migration speaks directly to the policy debate on neighborhood change, allows us to study cheaper segments of the market that may be under-represented in the Zillow™ data, and is the primary channel through which neighborhoods change (Freeman and Braconi 2004; McKinnish, Walsh, and White 2010; Brummet and Reed 2019).[7] In our near-near specification, we find that new construction decreases the average origin neighborhood income of in-migrants to the nearby area by about 2 percent. It also increases the share of in-migrants who are from very low-income neighborhoods by about three percentage points, suggesting that new buildings reduce costs in lower segments of the housing market, not just in the high-end units that are the most direct competitors of new buildings. Results are similar in the triple-difference specification and null in the near-far specification.[8]

Our exploration of where new market-rate rental apartments are built provides one explanation for why endogenous amenity or reputation effects may be small: new construction typically occurs after a neighborhood has already begun to change. Although we restrict to neighborhoods that are relatively low-income, those that receive new buildings are relatively high-education and experienced more income and education growth over the previous decade compared to neighborhoods that did not receive any new buildings. This suggests that rather than catalyzing demographic change in previously stable neighborhoods, new market-rate construction in low-income areas tends to follow neighborhood change, or gentrification. It

[7] We discuss out-migration later when presenting descriptive statistics but cannot study it directly due to data limitations described in Section 1.

[8] The migration results are generally noisier than the rent results. One possible explanation is that the Infutor data include both renters and homeowners, who may be differently affected by new rental construction. In addition, there may be a greater lag in the effect on migration outcomes, either because outcomes are actually slower to respond or because Infutor does not pick up a move immediately when it occurs.

may therefore provide little additional impetus for new amenities or signal that a neighborhood is now desirable. Instead, its primary effect appears to be to accommodate preexisting demand, diverting high-income households from nearby units and reducing rents.

These results matter for policy. Approving new housing in low-income areas is often contentious because of worries that new buildings will accelerate rent increases and gentrification. These local concerns also spill over into the regional policy debate and can stall large-scale housing reforms (Brey 2019; Dillon 2019). While these worries may be understandable given the dramatic changes occurring in many neighborhoods, our results suggest they are generally misplaced. Instead, policymakers should recognize that new market-rate housing has both local and regional benefits and should therefore be an important part of strategies to address the growing affordability crisis. In addition, our migration results suggest that strategies that encourage housing construction also foster more economically integrated neighborhoods, which could promote economic mobility for low-income residents (Chetty, Hendren, and Katz 2016; Chetty et al. 2018).

However, there are a few reasons for caution. First, our findings are specific to the large market-rate apartments and strong market cities that we study, and effects could differ for other types of housing or other areas if amenity effects depend on local context. Second, we are only able to follow outcomes for three years after building completion, though we provide evidence that longer-run effects are likely similar to our estimates. Finally, the actual implementation of reforms that increase housing supply requires changing complicated zoning and land-use regulations. Policymakers should keep in mind that the particulars of those changes could affect where housing is built—for example, in vacant lots or through demolition of existing affordable housing.

We contribute to a small but growing literature on the effects of new housing supply. Much of this work focuses on regional effects and uses models rather than quasi-experimental strategies. Favilukis, Mabille, and Van Nieuwerburgh (2019) and Nathanson (2019) develop calibrated spatial equilibrium models and find that new housing supply promotes regional affordability. Mast (2019) traces the sequence of moves induced by new construction to show that new construction loosens lower-income housing markets throughout a region. By contrast, Anenberg and Kung (2018) develop a neighborhood choice model and find small effects of new supply on prices. We complement these papers by focusing on the very local effects of new housing, which may differ from regional effects and play a large role in the current policy debate.

A few very recent papers also focus on the local effects of increased housing supply. Li (2019) finds that large new buildings in New York City lower nearby buildings' rental income even as they increase the number of nearby restaurants. By contrast, Singh (2019) studies

construction sparked by a 2006 property reform in New York City and finds that each new unit increases nearby buildings' rental income by over 2 percent (relative to a declining trend). We differ from these papers by studying many cities, using listing-level rent data, and examining migration outcomes.[9]

Finally, there is a large literature on land-use regulation, much of which focuses on the effects of regulatory barriers to housing construction.[10] It shows that restrictive regulations lead to higher rent and house price growth (Pollakowski and Wachter 1990; Quigley and Raphael 2004, 2005; Hilber and Vermeulen 2016), less migration into economically successful cities (Ganong and Shoag 2017), and city- and society-wide welfare losses (Turner et al. 2014; Hsieh and Moretti 2017; Bunten 2017; Parkhomenko 2017). Our research adds to this literature by suggesting that market-rate housing construction not only improves regional affordability but also neighborhood affordability.

The rest of this paper is organized as follows. Section 1 describes our building, rents, and mobility data and provides new descriptions of where market-rate housing is built. Section 2 describes our main empirical strategies. Section 3 presents the rent results, and Section 4 presents the migration results. Section 5 concludes.

1 Data and Summary Statistics

1.1 New Buildings Data

Data on large new rental apartment buildings are provided directly by Real Capital Analytics, a real estate market research firm that aims to track the universe of such buildings. Our starting sample includes 1,483 buildings with over 50 units that were completed in 2010–2019 in the following 11 central cities: Atlanta, Austin, Chicago, Denver, Los Angeles, New York City, Philadelphia, Portland, San Francisco, Seattle, and Washington, D.C.[11] The data include the name, year completed, number of units, and exact address of each building, which we geocode and match to 2010 census tracts (our definition of a neighborhood).

In order to focus on the types of buildings central to the policy debate and most likely to have large effects on amenities or neighborhood character, we impose several restrictions on

[9] Other work has studied adjacent questions about new development. Boustan et al. (2019) find that the introduction of condominiums did not increase central city income, education, or share white, which they attribute in part to the fact that most were built in areas that were already attractive to high-income households. Freemark (2019) finds that increasing zoned capacity in Chicago raised the value of affected land but does not study the effect on rents.

[10] See Gyourko and Molloy (2015) and Glaeser and Gyourko (2018) for recent reviews.

[11] We restrict to the borough of Brooklyn when studying New York City because this is the only borough included in our data use agreement with Zillow™.

our analysis sample.

1. *Large market-rate rental buildings:* We remove buildings that are income-restricted, senior or assisted living housing, and buildings flagged as student housing or located in tracts that are over 25 percent college students.[12] We also restrict to buildings with 50 units or more, as these are most contentious and represent the largest physical and direct demographic change to a neighborhood. Finally, we drop the small number of large condominium buildings in order to focus on the much more common rental buildings.

2. *Low-income central city neighborhoods:* We restrict to buildings constructed in initially low-income central city neighborhoods, which we define as census tracts that are in the principal city of their metropolitan area (defined as a Core-Based Statistical Area, or CBSA) and which have median household income below the CBSA median in the 2013–2017 American Community Survey.[13] We focus on low-income areas because they are the primary concern in the policy debate and are likely where amenity and reputation effects are largest.

3. *Sample years:* We restrict to buildings completed in 2015 and 2016 when analyzing effects on rents and 2014 and 2015 when analyzing effects on migration. This ensures that we have a sufficient number of years before and after treatment in our outcome data.

4. *Sufficient nearby listings:* We restrict to buildings with at least one observation (a listing for rental outcomes and a move for migration outcomes) in both the treatment and control areas in every year in the sample.[14] This restriction improves the precision of our estimates and helps to ensure that we are picking up buildings in preexisting neighborhoods where gentrification may be a concern, rather than large brownfield redevelopments.

5. *Isolated:* We include only buildings that had no other new buildings completed within 250 meters (the baseline treatment radius) between 2010 and their date of completion.

[12] We do so using a field provided with the data, which includes descriptions such as "subsidized" or "affordable." We do some spot checking to ensure that these flags are accurate.

[13] This is a common definition of "gentrifiable" used in the gentrification literature. We use this vintage of the ACS because it falls in the middle of our listing data. The final analysis sample changes by only one building if we instead use the 2006-2010 ACS.

[14] The definition of treatment and control areas changes depending on the specification. For the near-far approach, the treatment group is within 250 meters of a building, and the control is those between 250 and 600 meters. For the near-near approach, the treatment area is within 250 meters of early buildings, and the control is the same distance from later buildings.

This identifies relatively stable neighborhoods where a new building may represent a larger shock.

It is worth emphasizing that we are studying the effects of the buildings captured in this sample, not other sources of new housing such as single family homes, smaller apartment buildings, rehabs and renovations, or affordable housing developments. Moreover, our focus on isolated buildings precludes us from saying anything about the effect of building large clusters of buildings at the same time. However, our data suggest that clustering is limited. Of 1,483 total new buildings, 783 (53 percent) are isolated according to our definition, 336 (23 percent) are in a cluster of 2 total buildings, 178 (12 percent) are in a cluster of 3 total buildings, and the remaining 186 (13 percent) are in a cluster of 4 or more buildings. While large clusters of development such as Nationals Park in Washington, D.C., or Hudson Yards in New York City are interesting and salient, our data show they are relatively rare. Further, because these large clusters typically occur in industrial redevelopments rather than already established residential neighborhoods, they are less relevant to the policy debate. Finally, note that our sample does not drop buildings that successfully spur subsequent nearby construction, since we only require that buildings are isolated at the time of completion.

Table 1 shows building characteristics and location for four different subsamples. The first column, "All Incomes 2010–2019," restricts to large, market-rate buildings in central cities completed between 2010–2019 (restriction 1 in the above list). It gives an overall sense of the levels and types of construction occurring in our 11 cities. The second column further restricts to buildings built in low-income neighborhoods, which yields 823 buildings. The third column further restricts to those built in the years of our analysis sample (2014, 2015, or 2016), yielding 253. Finally, the last column adds the sufficient listings and isolation restrictions to yield our final rental analysis sample of 92 buildings. While the final sample is much smaller than the starting sample of 1,483, the table shows that most of the loss is simply from studying low-income neighborhoods in the years 2014 to 2016. The remaining decline is partly explained by the isolation restriction, which removes 106 buildings, and partly by the listings restriction, which tends to drop industrial redevelopments that are less relevant to the gentrification debate.

Despite the loss of observations, Table 1 shows that the different samples of low-income buildings are quite similar in both characteristics and location, giving us confidence that our analysis sample is relatively representative. The median number of units in our final sample is 121, though there is considerable variation: the 25th and 75th percentiles are 68 and 256. The nearby rent numbers show statistics for the Zillow™ listings within 600 meters of the new buildings (the baseline control radius).[15] The median nearby rent in our final sample is

[15]For each new building, we take the average of nearby listings. We then report the mean, median, and

$1,807 for all units and $1,514 for one-bedroom units.

We next consider what types of neighborhoods new developments are built in. Table 2 describes characteristics of central city neighborhoods (census tracts) that did and did not receive new construction between 2010 and 2019. The first two columns represent characteristics of neighborhoods of all income levels, with the first including only neighborhoods that did not receive construction and the second only those that did (weighted by the number of new buildings in each neighborhood), and columns 3 to 6 represent different groups of low-income neighborhoods. The rows show levels and growth rates for three sets of characteristics—income, education level, and rents—that are often used to characterize socioeconomic status and gentrification.[16]

The table reveals several interesting patterns that may help explain our main results. In both the all-income and low-income samples, neighborhoods with construction experienced much faster preperiod growth in college and income share and somewhat faster growth in rents. While the levels of income were similar across samples during the preperiod, the areas that received construction had higher college shares, which is often considered a leading indicator of gentrification. These results suggest that new buildings in low-income areas are typically built in neighborhoods that are already in early stages of gentrification. These areas may be appealing to developers because they are already attractive to the kinds of high-income residents who can pay the higher rents required for new construction.[17] In contrast, building market-rate units in low-income neighborhoods that are not gentrifying may be seen as riskier, as they may not be able to attract high-income residents.

Finally, despite our restrictions, the final analysis sample again appears to be relatively representative of buildings in low-income neighborhoods in general. In contrast, the neighborhoods that never receive construction are quite different, which cautions against generalizing our results to hypothetical buildings in the kinds of low-income, central city neighborhoods that developers currently rarely target.

other percentiles of the distribution of these building-level means.

[16]The table shows absolute levels and changes, but patterns are nearly identical when de-meaned within each CBSA. This and the similar distribution of building locations from Table 1 suggest that these patterns are unlikely to be driven by changes in sample composition.

[17]Another reason why developers target already-gentrifying neighborhoods could be that LIHTC investments can crowd out private development in neighborhoods that are lower income or in less-advanced stages of gentrification (Baum-Snow and Marion 2009; Eriksen and Rosenthal 2010). The Department of Housing and Urban Development (HUD) most heavily subsidizes LIHTC development in census tracts where at least 50 percent of residents are below 60 percent of the area median gross income or have a poverty rate of at least 25 percent. These tracts tend to be poorer than the neighborhoods we study here.

1.2 Rental Data

We pair the new building data provided by RCA with listing-level data on rental prices provided by Zillow™, which includes listings from all websites (Zillow™, Trulia, StreetEasy, and HotPads) in the Zillow™ Group during 2013 to 2018. For each listing, we observe price, location, date of listing, and the number of bedrooms and bathrooms.[18] The major restriction is that the sample includes only listings in buildings with 50 units or less. Because there are many such small buildings, this gives excellent geographic coverage, as demonstrated in Appendix Figure A.1, which shows the listings within 800 meters of a new building in Chicago in 2018. This dense coverage is useful for spatial identification strategies.

In total, we have rents for about 740,000 units within 800 meters of one of our new apartment buildings. Appendix Tables A.1 and A.2 show mean rent for all apartments and one-bedroom apartments in each CBSA-year. Rents generally increase over time and are higher in traditionally high-cost cities, but there is some variation in this pattern, possibly because of changes in which units are listed online. In our analysis, we always include CBSA $\times$ year fixed effects (or more granular effects) to absorb these compositional changes.

It is difficult to verify whether this sample of rents is representative, as Census Bureau statistics measure contracted rents, which could be quite different than the average listed rent at a given point in time. However, the mean rent for Zillow™ listings within 600 meters of a new building in our final sample, shown in Table 1, is $1,790, whereas the average ACS median rent in the tracts containing those buildings is only $1,165 (shown in Table 2). This suggests that Zillow™ rental listings skew toward higher-end units. Given that gentrifying areas often feature listings across a wide range of rents, this is important to bear in mind when interpreting our results. This later motivates us to use the migration data introduced below to directly study cheaper segments of the market. In addition, the differences between tract-level and nearby Zillow™ rents is consistent with developers targeting new construction to relatively higher-rent areas even within small neighborhoods. Similarly, Appendix Figure A.2 shows that rent (de-meaned at the CBSA-level) trends monotonically downward as distance to the nearest new building increases. This is exactly the endogeneity problem that our identification strategy aims to address.

1.3 Individual Migration Data

We use longitudinal individual address histories from Infutor Data Solutions to construct measures of migration. Infutor creates these data using numerous private and public record

[18] Because a small number of listings have extreme values, we winsorize the variable at the 1st and 99th percentile in all exercises.

sources—USPS change of addresses, county assessor records, magazine subscriptions, phonebooks, et cetera—and sells the data for use in targeted advertisements. Because addresses are intended to be used in direct mailing, they are high quality and reported at the unit level. Addresses are matched with an estimated arrival date, and there are limited demographics (age and gender) for each individual. The data closely match the census over-25 population at the tract level, with about 0.9 observations per census individual, and the coverage is similar across tract characteristics.[19] The Infutor data do miss some moves—the annual migration rate is 5.4 percent, versus 9.8 percent in the 2018 Current Population Survey—but this ratio appears to be uncorrelated with county characteristics.[20]

We use the address histories to measure the origin tract income of the individuals moving into the new buildings and into the nearby preexisting buildings.[21] These outcomes are useful for two main reasons. First, they are directly related to the policy debate about whether new buildings accelerate neighborhood demographic change or gentrification.[22] Second, they allow us to directly consider lower-income segments of the housing market by focusing on migration from low-income areas, whereas the Zillow™ data may skew toward the upper end of the housing market. A limitation of observing origin tract income rather than actual household income is that we cannot separate low- and high-income individuals within a tract, making it difficult to study low-income out-migration, which features prominently in the policy debate.

We observe about 1.9 million moves to an area within 800 meters of a new building between 2011 and 2017, including about 60,000 to a new building directly. Figure 1 shows trends in in-migration within 250 meters of the 2014–2015 buildings in our low-income analysis sample.[23] The triangle-marked line shows the number of migrants to the new building, the square-marked line shows migrants to the surrounding 250 meters excluding moves into the

[19] Appendix Figure A.3 plots the ratio of the Infutor and Census populations against tract characteristics. Additional validation is provided in Diamond, McQuade, and Qian (2019), which introduced the data to the academic literature.

[20] We assess this by computing the average annual migration rate at the county level in the Infutor data and comparing it to census estimates (which are not available below the county level). Appendix Figure A.4 plots the ratio of the two estimates against county characteristics and shows only slight correlations, suggesting that the moves we observe are close to randomly selected, Additionally, Phillips (2019) provides a detailed recent assessment of how well the Infutor data capture mobility among particularly disadvantaged populations, such individuals with very low incomes or initially living in public housing, and finds that they are good at capturing mobility and housing instability after events such as public housing demolitions or flooding from Hurricane Katrina.

[21] Infutor does not provide information on income for individuals themselves. Proxying for individual income with origin neighborhood income is similar to Mast (2019) and Diamond, McQuade, and Qian (2019).

[22] While much debate focuses on out-migration from gentrifying areas, the academic literature has reached a strong consensus that neighborhood change primarily occurs through changes in in-migration, making this the relevant outcome to study (Freeman 2004; McKinnish, Walsh, and White 2010; Ellen and O'Regan 2011; Ding, Hwang, and Divringi 2014; Brummet and Reed 2019).

[23] We exclude the small number of buildings with addresses that we cannot match to the Infutor data.

new building, and the hashed line is the sum of the two. Panel A counts migrants from tracts with above CBSA-median income, while Panel B tracks arrivals from tracts below two-thirds of the CBSA median.[24] Both cases show roughly the same story: the total number of arrivals of either type increases after building completion, but the change is totally driven by arrivals to the new building. Migration to the surrounding area changes little, though there is a slight decline in arrivals from high-income areas and a small uptick in arrivals from low-income areas. However, the new buildings clearly represent a substantial supply shock to the area: over 20 percent of high-income arrivals to the area in a given year are to the new building.

These results are inconsistent with the new buildings having a major impact on in-migration to the surrounding area, which is the primary way that neighborhoods change. However, because the public debate on gentrification largely concerns displacement of low-income households, we also consider net migration from low-income areas in Figure 2. Net migration does not appear to meaningfully change during the sample period. While this evidence is not causal, it is inconsistent with large displacement effects of new buildings.

2 Rental Empirical Strategy

2.1 Overview

To guide our empirical analysis, we first consider the hypothetical ideal experiment—building a new apartment building in a low-income neighborhood in a randomly assigned set of cities. We could then compare rental prices in the treatment and control group at different distances to the new building. Absent amenity or signaling effects, increased supply would likely cause prices to decrease close to the new building, and the effect would then fade out with distance (assuming that renter choice sets are at least somewhat spatially concentrated). But the shape of these other effects is much less clear, making it difficult to predict where the net effect on rents will be smaller or even positive. The ideal experiment would allow us to simply make comparisons at every distance, but in practice, different quasi-experimental strategies are better suited to detect different spatial patterns. We construct two empirical strategies to test for perhaps the most likely shapes of rent increases and consider some alternatives in extensions.

One possibility is that amenity effects are important very close to the new building and then fade out quickly. This makes sense for amenity changes like increased retail options, increased foot traffic, or aesthetic improvements like the replacement of a vacant lot. If these

[24] We exclude tracts with over 25 percent college students from the low-income migrants. We choose this low-income definition because it is a common threshold for income-restricted housing and because it provides a large enough sample to study.

effects are large, we could see rent increases very close to the new building that quickly fade out with distance. Our first empirical exercise—the "near-far" approach—checks for this pattern by comparing rents very close to a new building to rents slightly farther away. For identification, this approach relies on idiosyncratic variation in the availability of developable sites within small areas. Diamond and McQuade (2019) use similar logic to study new low-income housing developments, while Shoag and Veuger (forthcoming), Currie et al. (2013), and Autor et al. (2014) similarly use fine spatial variation in home locations to study the effect of localized shocks on prices.[25]

Another possibility is that amenity and reputation effects work at a broader level. For instance, a new building could signal that a relatively large neighborhood is on the upswing. This pattern would not be captured in our near-far strategy, as the control group would receive a similar boost as the treatment group. To get around this problem, we leverage variation in the timing of new developments, which varies due to the unpredictable length of the land assembly, permitting, and construction processes. This "near-near" approach compares the area around buildings completed in 2015–2016 to buildings completed in 2019 (after the end of our rent data). Because the control buildings can be in different neighborhoods than the treated, this allows us to detect broader effects of a new building. Finally, in addition to comparing the areas around the treatment and control buildings, we also construct a triple-difference approach that compares the near-far gap around the two sets of buildings.

2.2 Near-Far Approach

The near-far approach is a relatively standard "ring" difference-in-differences that compares listings within a treatment radius of the new building to listings within a larger control radius. The identification assumption is that in the absence of new construction, rents would have changed in parallel in the treatment and control group. The idea is that within a small area, developers are constrained in the sites that are available and have appropriate zoning, leading to hyper-local variation in the location of new construction that is not related to future price changes.

This identification assumption is more likely to hold with small ring radii that contain a relatively homogeneous area, while larger radii may introduce bias by including a control area that is dissimilar to the treatment area. However, the treatment group should also

[25] A review of the land-use literature also supports our identification assumptions. Ease of development can vary city by city, depending on zoning regulations and attitudes by city officials (Behroozi et al. 2001). Parcels may be small or irregularly shaped, and changes to their zoning may need to be made before development can proceed (Steinacker 2003). Developers preferring to circumvent residential opposition by building on formerly industrial or commercial sites may find that vacant land sites may be hard to assemble into sufficiently large parcels and be in poor physical conditions (Bowman and Pagano 2000; Brooks and Lutz 2016).

be large enough that it is substantially more affected by the new building than the control group. We set the baseline treatment radius to 250 meters (roughly one or two city blocks) and the control radius to 600 meters (slightly over a third of a mile, approximately an 8- to 10-minute walk).[26] The control radius is small enough to generally include listings that are part of the same neighborhood and likely on the same price trajectory. The treatment radius is large enough to include the listings in the immediate area of the building. We investigate the validity of our identification assumption by looking for pretrends in the event study graphs and consider alternate radii in robustness checks. Appendix Figure A.5 illustrates the treatment and control groups in Portland and Brooklyn.

To implement this strategy, we associate each listing i with the closest building b that was completed in 2015 or 2016 and define the treatment year t^* as that building's completion year. We keep listings with a building within 600 meters and index each listing's treatment status by r, which is equal to 1 if a listing is within 250 meters of its closest building and 0 otherwise. Because our rent data poorly identify when the same unit is listed at multiple points in time, we cannot include unit fixed effects and must instead use fixed effects at the level of the nearest building $\times$ treatment status. We also use the nearest building to define time fixed effects that control for time-varying shocks at a very local level. This yields the following specification for rent in listing i in year t:

$$log(rent_{it}) = \alpha_{bt} + \kappa_{br} + \sum_{k=-3}^{3} [\beta_k * \mathbb{1}_{it}(t - t^* = k, r = 1)] + \gamma * X_{it} + \epsilon_{it}. \tag{1}$$

X_{it} consists of dummies for the number of bedrooms and bathrooms, which helps to control for changes in the composition of listings over time. We weight each building-year equally to account for different densities of nearby listings and cluster standard errors at the level of the nearest new building. We also estimate a standard difference-in-differences (DiD) to obtain an average effect:

$$log(rent_{it}) = \alpha_{bt} + \kappa_{br} + \beta * \mathbb{1}_{it}(t \geq t^*, r = 1) + \gamma * X_{it} + \epsilon_{it}. \tag{2}$$

2.3 Near-Near Approach

The previous approach tests for highly localized spillovers of the new buildings but, because the treatment and control group are very close together, cannot identify more spatially diffuse

[26] Time estimates are from Google Maps and vary depending on how directly the road network connects two sites. While the treatment area is small, both amenity and supply effects of new housing have been found to decay quickly with distance. Schwartz et al. (2006) and Rossi-Hansberg et al. (2010) find that positive housing externalities from urban revitalization programs substantially decay within 250 meters. Campbell, Giglio, and Pathak (2011) and Hartley (2014) find that the negative effect of a foreclosure decays even faster.

effects. To allow for this type of spillover, our near-near strategy compares the area around 2015–2016 buildings to a set of similar locations: the area around 2019 buildings. Because the treatment and control group are no longer necessarily in the same neighborhood, even broader spillovers from the treatment buildings should not affect the control listings. In addition, this specification buttresses our near-far analysis by using an entirely different source of variation.

The identification assumption is that rents would have changed in parallel near the two sets of buildings in the absence of construction. The rationale is that, because both treatment and control locations received new construction in a four-year period, they share characteristics that make them attractive to developers. Moreover, variation in timing of completion could be largely due to to the idiosyncratic nature of the land assembly, building permitting, and construction processes.

Table 3 helps evaluate this assumption by showing differences between neighborhoods near treatment and control buildings. We regress characteristics of each building's tract on CBSA fixed effects and a treatment dummy. While none of the differences are statistically significant at the 5 percent level, the point estimates suggest that the treatment neighborhoods are somewhat wealthier, whiter, and more expensive. This likely occurs because the earlier buildings went into areas that began to gentrify at least somewhat earlier. However, the control group scores higher on some indicators: it has a higher home ownership rate and a lower vacancy rate. On the whole, the treatment and control neighborhoods appear to be similar, and we further probe the identification assumption by studying pretrends in the event study results.

To implement this approach, we now associate each listing with its nearest 2015, 2016, or 2019 building and keep only those with a building within 250 meters. Because there is no variation in the treatment variable within listings that share the same nearest building, we coarsen our time fixed effects from nearest building $\times$ year to CBSA $\times$ year and similarly coarsen our nearest building $\times$ treatment status fixed effects to the nearest building level.[27] We index a listing's treatment status by c, which is equal to 1 if a listing's nearest building was completed in 2015 or 2016 and 0 otherwise.

This yields the following event study specification for listing i in CBSA m in year t:

$$log(rent_{it}) = \alpha_{mt} + \kappa_b + \sum_{k=-3}^{3} [\beta_k * \mathbb{1}_{it}(t - t^* = k, c = 1)] + \gamma * X_{it} + \epsilon_{it}. \tag{3}$$

We weight observations and cluster standard errors as in the near-far specification, and we

[27] Because San Francisco and Washington, D.C., do not have 2019 buildings that meet our sample criteria, we drop these CBSAs from the sample for the near-near specification.

also estimate the standard DiD:

$$log(rent_{it}) = \alpha_{mt} + \kappa_b + \beta * \mathbb{1}_{it}(t \geq t^*, c = 1) + \gamma * X_{it} + \epsilon_{it}. \tag{4}$$

2.4 Triple-Difference Approach

Finally, our first two specifications can be combined naturally into a triple-difference approach. This can be interpreted as relaxing the identification assumption in either the near-far or near-near specification. Relative to the near-far, it effectively uses the near-far gap around 2019 buildings to control for time-varying omitted variables that may affect the same gap near 2015–2016 buildings. This changes the identification assumption to: in the absence of construction, the near-far rent differential would have changed in parallel in the 2015–2016 and 2019 areas. This allows for developers to have some discretion in site selection, even within our small radius, as long as that discretion is the same in both 2015–2016 and 2019.

Alternatively, the triple difference can be viewed as improving on the near-near approach by comparing the difference between the outer rings of the treatment and control buildings to the difference between the inner rings of the two sets of buildings. This relaxes the identification assumption to: in the absence of construction, the treatment-control rent differential would have changed in parallel in the inner and outer rings. This allows for differences between the treatment and control inner rings (such as the treatment group gentrifying earlier), so long as the difference is the same in the treatment and control outer rings.

To implement the triple-difference specification, we must first expand the sample to include all listings within 600 meters of either a treatment or control building. Noting that r indexes whether an observation is within 250 meters of its closest building (treated in the near-far) and c indexes whether an observation is near a 2015–2016 building (treated in the near-near), the specification is:

$$log(rent_{it}) = \alpha_{bt} + \kappa_{br} + \nu_{rt} + \beta * \mathbb{1}_{it}(t \geq t^*, c = 1, r = 1) + \gamma * X_{it} + \epsilon_{it}, \tag{5}$$

where ν_{rt} is a fixed effect for the inner ring in each year. We use the same weighting and clustering schemes as in the previous specifications.

3 Rent Results

3.1 Main Results

Results for the near-far specification are in Figure 3 and Table 4, Column 1. The event study specification in Figure 3 shows that coefficients (normalized to the year before completion) hover around zero during the preperiod and then sharply drop to roughly -0.06 after building completion. The effect remains very similar for the three subsequent years that we can track in our sample. Estimating the model using an after-treatment dummy (Table 4, Column 1) shows that new buildings decrease nearby rents by 4.9 percent (S.E.=0.021, p=0.023). These results suggest that hyper-local positive spillovers do not cause rents to rise in the immediate area of the building.

Results for the near-near specification, shown in Figure 3 and Table 4, Column 2, also suggest that new buildings decrease rents. The event study coefficients are similar to those from the near-far specification: approximately zero before falling to about -0.07 following building completion. The after-treatment estimate from a standard DiD is -0.062 (S.E.=0.037, p=0.096).

Finally, the triple-difference estimate is in Table 4, Column 3. It is again quite similar to the near-far and near-near estimates: -0.071 (S.E.=0.033, p=0.037). In dollar terms, these point estimates translate to between a \$100 and \$159 decrease in listing rent, which is about \$1 per unit in the new building on average. We note that if developers are less constrained in their ability to select locations than our specifications assume, this should lead to a positively selected treatment group, likely biasing our results toward 0.

The baseline estimates use a relatively small treatment radius of 250 meters, which we extend to 400 meters in Appendix Table A.3.[28] The near-near estimate decreases to a statistically insignificant -2.8 percent (S.E.=0.032, p=0.39), and the near-far estimate similarly shrinks. However, the event study coefficients underlying these specifications (Appendix Figures A.6 and A.7) show that there is a strong positive pretrend in the treated area. The pretrend illustrates the hyper-local nature of our identification assumptions: when we push the control group slightly further away from the new building, it no longer provides an ideal comparison. Nonetheless, the evidence from the 400-meter exercises still strongly points to a negative rent effect of the new building. The positive pretrend in the event studies sharply reverses in the treatment year, and our triple-difference specification, which accounts for the positive pretrend, yields a negative effect of 6.8 percent (S.E.=0.041, p=0.097).

Before moving to extensions and robustness, we discuss two caveats that we cannot explore

[28] We also extend the isolation restriction on new buildings to 400 meters to match the treatment radius.

with our rent data. First, significant heterogeneity across building and neighborhood types may underlie our average results. Amenity effects likely vary highly depending on the context, and there may be important tipping points or nonlinearities. For example, buildings may have a different effect if they replace a vacant lot versus an existing building, if they offer some particularly attractive amenity, or depending on the current national or local economic climate. Our sample size prevents us from exploring this heterogeneity in detail, and we note that, even given a larger sample, we would not be able to estimate the effect of new buildings in the types of neighborhoods where developers never choose to build. Second, the data from Zillow™ tend to skew towards the higher end of the rent distribution, which makes it difficult to assess whether the average effects we estimate might be different for the initially lower-rent units that are most relevant to the policy debate. While we have no other source of listing-level rent data to study this directly, our migration data allow us to study the lower end of the housing market by examining migration from low-income areas.

3.2 Congestion Effects and Other Spatial Patterns

Our near-far and near-near strategies aim to detect rent effects that center around new buildings, but these effects could take many shapes, depending on the interaction and size of amenity, disamenity, and supply effects. For example, new buildings could improve consumption amenities within a ten minute walk of the new building, but also generate congestion disamenities by increasing traffic or blocking views on the building's block. This would yield a "doughnut-shaped" rent effect in which rents decline nearby but increase slightly farther away. To test this idea, we repeat the near-near strategy but compare the 250- to 600-meter bands in the treatment and control groups. Appendix Table A.4, shows statistically insignificant effects of 2.8 percent (S.E. = 0.026, p=0.27) and 1.6 percent (S.E. = 0.021, p=0.46) for the areas between 250–600 meters and 250–800 meters, respectively.[29]

While congestion likely follows a predictable pattern, cities are complex systems in which a new building could generate other nonlinear or oddly shaped amenity effects. We evaluate these possibilities more generally using Diamond and McQuade's (2019) empirical derivative approach, which produces continuous estimates of treatment effects on rents at various distances to the new buildings. The primary drawback is that it only accommodates buildings that had no other building within the studied area during the entire sample period, which reduces our sample size and results in a potentially selected set of buildings.[30] Appendix Figure A.8 shows the estimated treatment effect at distances within 600 meters of the new

[29] While these point estimates do not allow us to rule out positive effects, we note that point estimates are negative when we do a similar exercise with migration outcomes in the next section.

[30] We discuss implementation in more detail in the Appendix.

building, both before and after its completion. They are analogous to the near-far event study coefficients and show a very similar pattern. Across distances, rents change in parallel prior to the building's completion. After completion, rents decrease much more sharply closer to the building, especially in the area within 250 meters that represents our treatment radius in the near-far specification.[31] Beyond 250 meters, the gradient rises toward zero as we approach the 600-meter limit. This simple pattern give us confidence that our ring approach is not netting out any nonlinearities that would undermine our main results.

3.3 Robustness

Our first set of robustness checks explores the sensitivity of the baseline results to alternate assumptions. Appendix Figure A.9 shows the near-far event study under four different restrictions—dropping majority white tracts, reducing the income threshold by $10,000, requiring that buildings also be in below-median college share neighborhoods, and dropping listings in New York.[32] In all cases, the patterns appear similar to the baseline, and the DiD estimates are again between −5 percent and −7 percent (Appendix Table A.5). We repeat these exercises for the near-near specification and find that results are similar to the baseline, including a larger negative effect when we restrict to low-college or highly nonwhite areas (Appendix Figure A.10 and Appendix Table A.6).

Next, we try to examine longer-run effects and longer pretrends, although we cannot do so with our main analysis sample. First, we repeat the near-far specification using 2010–2013 buildings, which allows us to include five postconstruction years (Appendix Figure A.11), and find no evidence of long-run changes in the treatment effect.[33] We then repeat the same specification using only 2019 buildings, which allows us to study pretrends over the six years before completion. We see no evidence of time trends and no discernible effect of project approvals or announcements that precede building completion, although our small number of buildings (23) leads to some noisy estimates (Appendix Figure A.12). We also study longer-run pretrends using the empirical derivative approach (Appendix Figure A.13)

[31] The estimated effect near the construction site is almost −40 percent, which is much higher than the rest of our results. Beyond the selection bias introduced by the more stringent isolation criteria discussed above, it could also be due to the challenges inherent in adopting an algorithm designed to estimate effects on housing transactions to estimating effects on apartment listings. Housing transactions are rarer, and neighboring houses in central cities tend to be more similar than apartment units. The fact that apartment units may be more heterogeneous within a given area likely makes spatial differencing more difficult.

[32] The number of listings and average rent in New York change sharply in 2017, as shown in Appendix Table A.1. Although we include CBSA × year or nearest building × year fixed effects in all specifications, dropping these observation provides another check that this change does not drive our results.

[33] We include four years of buildings because, in the wake of the Great Recession, there was little construction in the 2010–2012 time period. We do not do this exercise for the near-near specification because the 2019 buildings are a substantially worse control group for 2010 buildings than for 2015 and 2016 buildings.

and find that this more detailed price surface tells a similar story. While these exercises do not use our main analysis buildings and we cannot study both longer-run pretrends and longer-run effects within the same specification, the results nevertheless provide suggestive evidence that our main estimates would not drastically change if we were able to extend our time window to five years.

Finally, we remove the restriction on neighborhood income and repeat our main specifications. Figure A.14 shows the near-far event studies for both 250- and 400-meter radii, and Figure A.15 shows the same for the near-near specification. The near-far results show no consistent pattern, and the near-near results show noisy patterns that are, at most, somewhat suggestive of a negative effect of new buildings. Table A.7 shows DiD estimates with the 250-meter radius, which are close to zero and not statistically significant. Table A.8 shows the estimates for the 400-meter radius, which are similar except for the triple difference, which is statistically significant and negative. In some ways, the smaller estimates appear counterintuitive, as new buildings likely represent less of a change for richer neighborhoods, suggesting that the positive amenity effect should be smaller. However, the elasticity of demand for rich neighborhoods may be higher if they are broadly liked by high-income households and gentrifying or low-income neighborhoods appeal to only a smaller subset. This would reduce the price effect of increasing supply. In addition, the near-far strategy in particular may be less valid in high-income areas if well-organized NIMBY opposition is better able to shunt new buildings to specific sites that are unattractive or dissimilar to their surroundings.

4 Migration Results

4.1 Migration Empirical Strategy

Our empirical strategy is extremely similar to the rental analysis described in Section 2, with the following differences. First, instead of a rental listing, an observation is now an in-migrant to an address near a new building. Second, instead of using price as the dependent variable, we use median household income in the in-migrant's origin tract. Third, because the Infutor sample ends in 2017, we use 2014–2015 buildings as the primary treatment group and 2018 buildings as the control. Finally, we cannot include controls for bedroom and bathroom counts, which we do not observe in the migration data.

With these changes, we estimate event studies and DiDs for the near-far approach (Equations 1 and 2), near-near approach (Equation 3 and 4), and the triple difference (Equation 5). Because our goal is to assess how new buildings affect the surrounding

neighborhood, we always exclude from the sample migrants to the new building itself.[34]

4.2 Migration Results

We begin with results from the near-near specification. Event study coefficients for log origin income are shown in Figure 5, Panel A. There is a parallel trend in the preperiod followed by a sharp decrease of approximately 3–5 percent in the three years following building completion. The drop appears to have a 1–2 year delay after building completion, which could occur either because migration is slower to respond to a supply shock than listed rents or because move dates in the Infutor data have a lag, as the address change may not be detected immediately upon moving. Because of this lag, we include the year of building completion as part of the preperiod and drop the year after treatment from the difference-in-differences estimation. As shown in Table 5, Panel A, the average effect of a new building on log origin income is -0.03 (S.E. = 0.017, p=0.085).

The result for origin income provides some corroboration of the rent results using an alternative data source, but it does not speak directly to effects on the high- or low-income segments of the housing market. To measure these effects, we repeat the specification using indicators for whether an in-migrant is from a high-income tract (above CBSA median) or low-income tract (below two-thirds of the CBSA median) as the dependent variable. Event study results are shown in Figure 5, Panel B. For low-income arrivals, there is a parallel pretrend that breaks sharply upward one year after treatment. The pattern is nearly the opposite for high-income arrivals, although the pretrend is noisier and slightly upward. The average effects, shown in Panels B and C of Table 5, are 0.034 (S.E. = 0.017, p=0.059) for low-income arrivals and -0.023 (S.E. = 0.021, p=0.295) for high-income arrivals. For low-income arrivals, this represents an approximately 15 percent increase to the mean in the year before completion. These results suggest that the effect of new buildings on rents is not driven entirely by the high-end listings that are their closest competitors. Instead, the increased migration from low-income areas is consistent with rent decreases among relatively cheaper apartments. Apart from their implications for rent, these results suggest that construction allows more low-income households to move to or remain in the sample neighborhoods, most of which appear to be gentrifying.

By contrast, the near-far results are not very informative. The event study results in Figure 6 show no consistent pattern: across all three outcomes, the years-to-treatment coefficients are small and do not show any strong trend. Similarly, the difference-in-differences estimates in the first column of Table 5 all have t-statistics well below 1. While these null

[34]For the small number of buildings with addresses that we cannot match to the Infutor data, we drop all migrants to addresses within 20 meters of the building.

results may be surprising given the statistically significant near-near results, there are a number of reasons that this combination of dependent variable and specification may have the least power. First, new buildings are rentals, but our migration outcomes include both renters and homeowners, whom we cannot separate. If homeowner location choices respond less to shocks to the rental market, this may push our estimates toward zero. Second, the near-far strategy generally has less power because some effects of the new building may affect the control group. Last, in the near-far migration specification in particular, variation in the percentage of renters and homeowners at different distances to the new building could lead to added noise that is not present in the near-far rent specification. This may occur if new rental buildings tend to be constructed in areas that are already zoned for relatively high-density housing.

Finally, the third column of Table 5 shows migration estimates from the triple-difference specification. The magnitudes are quite similar to those of the near-near estimates, with a negative effect on origin income that is significant at the 10 percent level and a positive effect on low-income arrivals that is significant at the 5 percent level.

4.3 Migration Robustness

We repeat a number of the extensions and robustness checks from the rental analysis, focusing on the near-near specification that yielded the most compelling results. First, Appendix Figures A.16 and A.17 show event study results for log origin income and high- and low-income arrivals under a number of restrictions to the sample: reducing the income threshold, restricting to less-white or lower-education neighborhoods, and dropping New York City. All show similar patterns to the baseline results. In addition, we extend the treatment ring to 400 meters. The event studies in Appendix Figure A.18 are loosely similar to the 250-meter radius, but the patterns are less clear.

Next, again motivated by the possibility of nearby congestion effects creating doughnut-shaped net amenity spillovers, we run the near-near analysis using arrivals to the 250–600 and 250–800 meter bands in the treatment and control groups. As shown in Appendix Table A.9, we find small negative point estimates on origin income in both samples, but neither is statistically significant. Like the analogous rent results, this provides no evidence that new buildings increase rents in this alternative pattern.

5 Conclusion

Prior research has shown that new market-rate housing construction improves regional housing affordability, but there is little evidence on how it affects the immediately surrounding neighborhood. This leaves open questions of incidence: who are the winners and losers of new housing construction? The housing policy debate is heavily influenced by the belief that new construction, especially in low-income areas, makes local incumbents worse off by attracting high-income households and high-end amenities, which in turn raises nearby rents and accelerates demographic change (Been, Ellen, and O'Regan 2019; Hankinson 2018). This belief has played a large role in the defeat of major housing reforms such as California's Senate Bill 50 and is frequently invoked at community meetings that influence the approval of individual housing development proposals.

We attempt to fill this gap in the literature using quasi-experimental methods and address-level microdata, focusing on the low-income areas that are most relevant to the policy debate. We find that the concerns that motivate opposition to new market-rate housing are mostly unfounded. While there is a strong observed correlation between new construction, rising rents, and demographic change, this is because new buildings are typically constructed in areas that are already changing. When these new buildings are completed, they actually slow rent increases and demographic change in the nearby area. The average new building lowers nearby rents by 5 to 7 percent relative to trend, translating into a savings of $100–$159 per month. Results are consistent across a number of specifications, and remaining bias from unobservables driving the selection of building sites likely pushes our estimates toward zero. In addition, we find that new buildings increase low-income in-migration, implying that this improved affordability can foster more integrated, economically diverse neighborhoods that may provide low-income residents with greater economic mobility (Chetty, Hendren, and Katz 2016; Chetty et al. 2018).

The mechanism underlying these results appears to be a simple story of supply and demand. If high-income households like a particular neighborhood, preventing the construction of new housing in those neighborhoods does not prevent them from moving to that neighborhood. Instead, it simply leads them to outbid lower-income households for whatever housing is already available in that neighborhood. This raises rents for everyone and lowers the ability of low-income residents to stay in or move to the area. By contrast, if new housing is built, many high-income households will choose this option instead of a nearby existing unit, reducing rent and out-migration pressures in the area. The new building could theoretically change local amenities or reputation by enough to instead increase demand and raise rents for nearby units, but our findings suggest this is not the case. Our descriptive finding that new buildings

go into areas that are already changing could explain why it is not the case: when an area's amenities and reputation are already changing for other reasons, the marginal effect of a new building on amenities and reputation may be small.

Increasing housing supply should therefore be an important part of any solution to the present affordability crisis. One way to do so is to relax regulations that make it difficult to build in many cities (Gyourko and Molloy 2015; Ganong and Shoag 2017; Glaeser and Gyourko 2018), and another is to directly incentivize localities to increase housing production. However, we note several caveats to our findings. First, we estimate an average effect that may disguise heterogeneity across building and neighborhood types, as amenity and reputation effects are likely highly context-dependent. Second, our sample consists of areas where developers actually chose to build. While this is most relevant to the debate, as these are likely the neighborhoods that would receive new construction if supply were increased, effects may be different in other types of neighborhoods. For example, developers do not build market-rate units in very low-income areas with high vacancy rates, so our results do not speak to what would happen if they did. Third, relaxing land-use regulation is, in practice, quite complicated. The particulars of a reform could matter both for how much supply is actually added and, depending on the incentives built in (such as encouraging redevelopment of the existing housing stock versus vacant land), the local effect of that new supply. These caveats point to important areas for future research.

References

Albouy, D., Ehrlich, G., and Liu, Y. (2016). Housing demand, cost-of-living inequality, and the affordability crisis. NBER Working Paper 22816.

Anenberg, E. and Kung, E. (2018). Can more housing supply solve the affordability crisis? Evidence from a neighborhood choice model. *Regional Science and Urban Economics*.

Baum-Snow, N. and Marion, J. (2009). The effects of low income housing tax credit developments on neighborhoods. *Journal of Public Economics*, 93(5):654–666.

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

Behroozi, C., Eakin, T., Suchman, D. R., and Todd, J. (2001). Urban infill housing: Myth and fact. Urban Land Institute.

Boustan, L. P., Margo, R. A., Miller, M. M., Reeves, J. M., and Steil, J. P. (2019). Does condominium development lead to gentrification? NBER Working Paper 26170.

Bowman, A. O. and Pagano, M. A. (2000). Transforming America's cities: Policies and conditions of vacant land. *Urban Affairs Review*, 35(4):559–581.

Brey, J. (2019). Why some tenant groups are opposing California's density bill. Next City.

Brooks, L. and Lutz, B. (2016). From today's city to tomorrow's city: An empirical investigation of urban land assembly. *American Economic Journal: Economic Policy*, 8(3):69–105.

Brummet, Q. and Reed, D. (2019). The effects of gentrification on the well-being and opportunity of original resident adults and children. Federal Reserve Bank of Philadelphia Working Paper 19-30.

Bunten, D. (2017). Is the rent too high? Aggregate implications of local land-use regulation. Board of Governors of the Federal Reserve System Finance and Economics Discussion Series.

Campbell, J. Y., Giglio, S., and Pathak, P. (2011). Forced sales and house prices. *American Economic Review*, 101(5):2108–31.

Chetty, R., Friedman, J. N., Hendren, N., Jones, M. R., and Porter, S. R. (2018). The opportunity atlas: Mapping the childhood roots of social mobility. NBER Working Paper 25147.

Chetty, R., Hendren, N., and Katz, L. F. (2016). The effects of exposure to better neighborhoods on children: New evidence from the Moving to Opportunity Experiment. *American Economic Review*, 106(4):855–902.

Diamond, R. (2016). The determinants and welfare implications of U.S. workers' diverging location choices by skill: 1980-2000. *American Economic Review*, 106(3):479–524.

Diamond, R. (2017). Housing supply elasticity and rent extraction by state and local governments. *American Economic Journal: Economic Policy*, 9(1):74–111.

Diamond, R. and McQuade, T. (2019). Who wants affordable housing in their backyard? An equilibrium analysis of low income property development. *Journal of Political Economy*, 127(3).

Diamond, R., McQuade, T., and Qian, F. (2019). The effects of rent control expansion on tenants, landlords, and inequality: Evidence from San Francisco. *American Economic Review*, 109(9):3365–94.

Dillion, L. (2019). California bill to add housing in single-family home neighborhoods blocked by lawmakers. *Regional Science and Urban Economics*.

Ding, L., Hwang, J., and Divringi, E. (2016). Gentrification and residential mobility in Philadelphia. *Regional Science and Urban Economics*, 61:38–51.

Ellen, I. G. and O'Regan, K. M. (2011). Gentrification: Perspectives of economists and planners. *The Oxford Handbook of Urban Economics and Planning*.

Eriksen, M. D. and Rosenthal, S. S. (2010). Crowd out effects of place-based subsidized rental housing: New evidence from the LIHTC program. *Journal of Public Economics*, 94(11):953 – 966.

Favilukis, J., Mabille, P., and Van Nieuwerburgh, S. (2019). Affordable housing and city welfare. NBER Working Paper 25906.

Freeman, L. (2005). Displacement or succession? Residential mobility in gentrifying neighborhoods. *Urban Affairs Review*, 40(4):463–491.

Freeman, L. and Braconi, F. (2004). Gentrification and displacement New York City in the 1990s. *Journal of the American Planning Association*, 70(1):39–52.

Freemark, Y. (2019). Upzoning Chicago: Impacts of a zoning reform on property values and housing construction. *Urban Affairs Review*.

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

Guerrieri, V., Hartley, D., and Hurst, E. (2013). Endogenous gentrification and housing price dynamics. *Journal of Public Economics*, 100:45–60.

Gyourko, J., Mayer, C., and Sinai, T. (2013). Superstar cities. *American Economic Journal: Economic Policy*, 5(4):167–99.

Gyourko, J. and Molloy, R. (2015). Regulation and housing supply. In Duranton, G., Henderson, J. V., and Strange, W. C., editors, *Handbook of Regional and Urban Economics*, volume 5 of *Handbook of Regional and Urban Economics*, pages 1289–1337. Elsevier.

Hankinson, M. (2018). When do renters behave like homeowners? High rent, price anxiety, and nimbyism. *American Political Science Review*, 112(3):473–493.

Hartley, D. (2014). The effect of foreclosures on nearby housing prices: Supply or dis-amenity? *Regional Science and Urban Economics*, 49:108–117.

Hilber, C. A. L. and Vermeulen, W. (2016). The impact of supply constraints on house prices in England. *Economic Journal*, 126(591):358–405.

Hsieh, C.-T. and Moretti, E. (2019). Housing constraints and spatial misallocation. *American Economic Journal: Macroeconomics*, 11(2):1–39.

Li, X. (2019). Do new housing units in your backyard raise your rents? Working Paper.

Mast, E. (2019). The effect of new market-rate housing construction on the low-income housing market. Working Paper.

McKinnish, T., Walsh, R., and White, T. K. (2010). Who gentrifies low-income neighborhoods? *Journal of Urban Economics*, 67(2):180–193.

Nathanson, C. (2019). Trickle-down housing economics. Working Paper.

Parkhomenko, A. (2018). The rise of housing supply regulation in the U.S.: Local causes and aggregate implications. (275).

Phillips, D. (2019). Measuring housing stability with consumer reference data. *Demography*, upcoming.

Pollakowski, H. O. and Wachter, S. M. (1990). The effects of land-use constraints on housing prices. *Land Economics*, 66(3):315–324.

Quigley, J. M. and Raphael, S. (2004). Is housing unaffordable? Why isn't it more affordable? *Journal of Economic Perspectives*, 18(1):191–214.

Quigley, J. M. and Rosenthal, L. A. (2005). The effects of land use regulation on the price of housing: What do we know? What can we learn? *Cityscape*, 8(1):69–137.

Rosenthal, S. S. (2014). Are private markets and filtering a viable source of low-income housing? Estimates from a repeat income model. *American Economic Review*, 104(2):687–706.

Rossi-Hansberg, E., Sarte, P., and Owens, R. (2010). Housing externalities. *Journal of Political Economy*, 118(3):485–535.

Schwartz, A., Ellen, I., Voicu, I., and Schill, M. (2006). The external effects of place-based subsidized housing. *Regional Science and Urban Economics*, 36(6):679–707.

Singh, D. (2019). Do property tax incentives for new construction spur gentrification? Evidence from New York City. Working Paper.

Steinacker, A. (2003). Infill development and affordable housing: Patterns from 1996 to 2000. *Urban Affairs Review*, 38(4):492–509.

Su, Y. (2018). The rising value of time and the origin of urban gentrification. Working Paper.

Turner, M. A., Haughwout, A., and van der Klaauw, W. (2014). Land use regulation and welfare. *Econometrica*, 82(4):1341–1403.

Figure 1: In-Migration to Areas around New Buildings

Panel A: Origin Income > CBSA Median

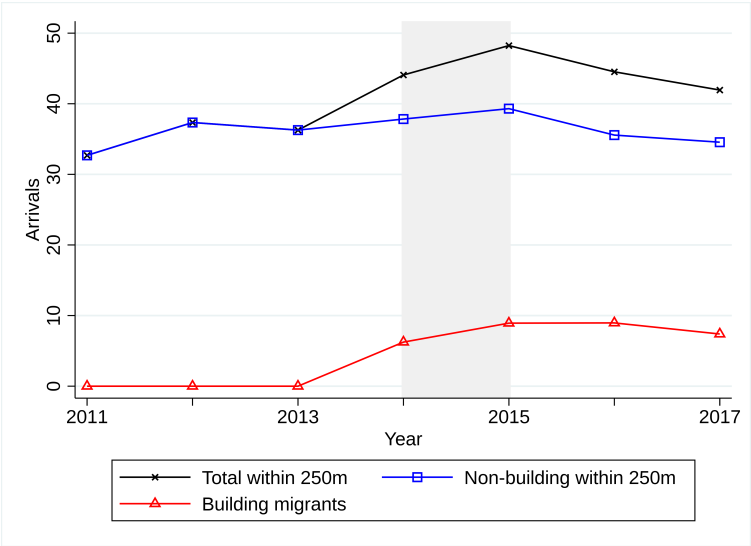

Panel B: Origin Income <2/3 CBSA median

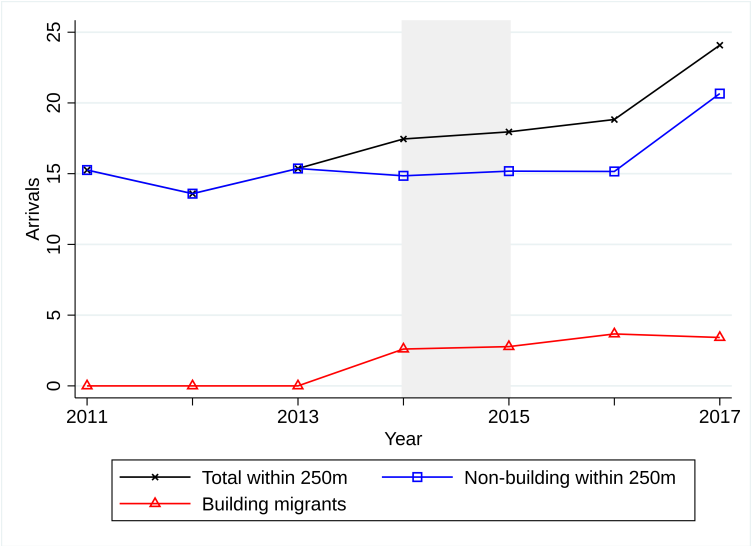

Note: This figure shows trends in the number of in-migrants to the area within 250 meters of new buildings. Panel A counts only in-migrants whose previous address is in a census tract with income above the CBSA median, while Panel B restricts to below two-thirds of the CBSA median income. Nonbuilding migrants are those arriving to the area within 250 meters but not the new building, building migrants are arrivals to the new building itself, and total migrants is the sum. The sample includes 2011–2017 moves to the area within 250 meters of the 2014–2015 buildings in our final analysis sample. Income in the origin tract is defined as median household income in the 2013–2017 ACS, and tracts with over 20 percent college students are excluded from the low-income definition. Migration totals in each year are normalized to account for differences in aggregate Infutor coverage across years, and we drop the small number of new buildings with addresses that we cannot match to the Infutor data.

Figure 2: Net Migration between Low-Income Tracts and New Building Areas

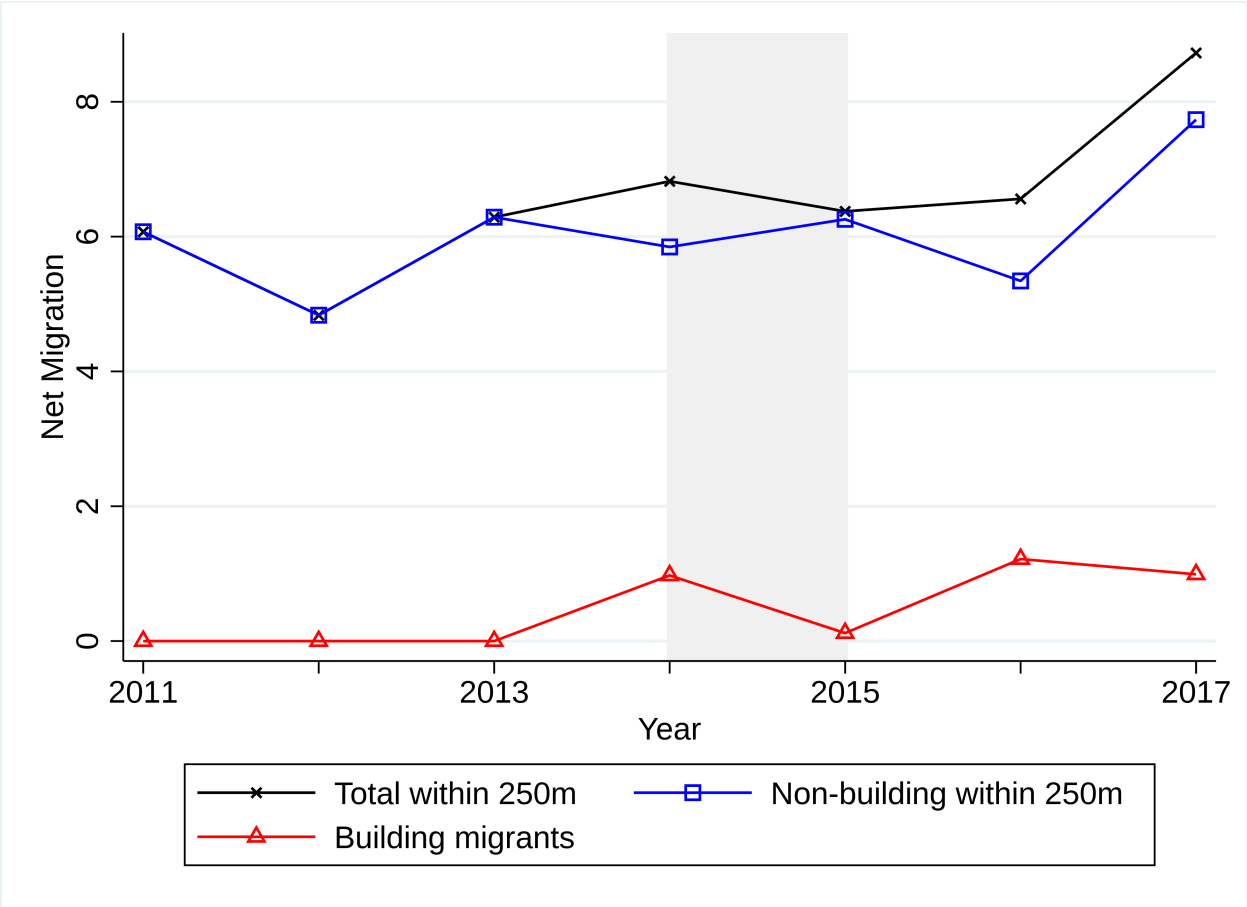

Note: This figure shows trends in net migration (arrivals - departures) between tracts with income below two-thirds of the CBSA median and the area within 250 meters of new buildings. Nonbuilding migrants are those arriving to the area within 250 meters but not the new building, building migrants are arrivals to the new building itself, and total migrants is the sum. The sample includes 2011–2017 moves and the area within 250 meters of the 2014–2015 buildings in our final analysis sample. Income in the origin tract is defined as median household income in the 2013–2017 ACS, and tracts with over 20 percent college students are excluded. Migration totals in each year are normalized to account for differences in aggregate Infutor coverage across years, and we drop the small number of buildings with addresses that we cannot match to the Infutor data.

Figure 3: Near Versus Far Event Study for Rent Outcome

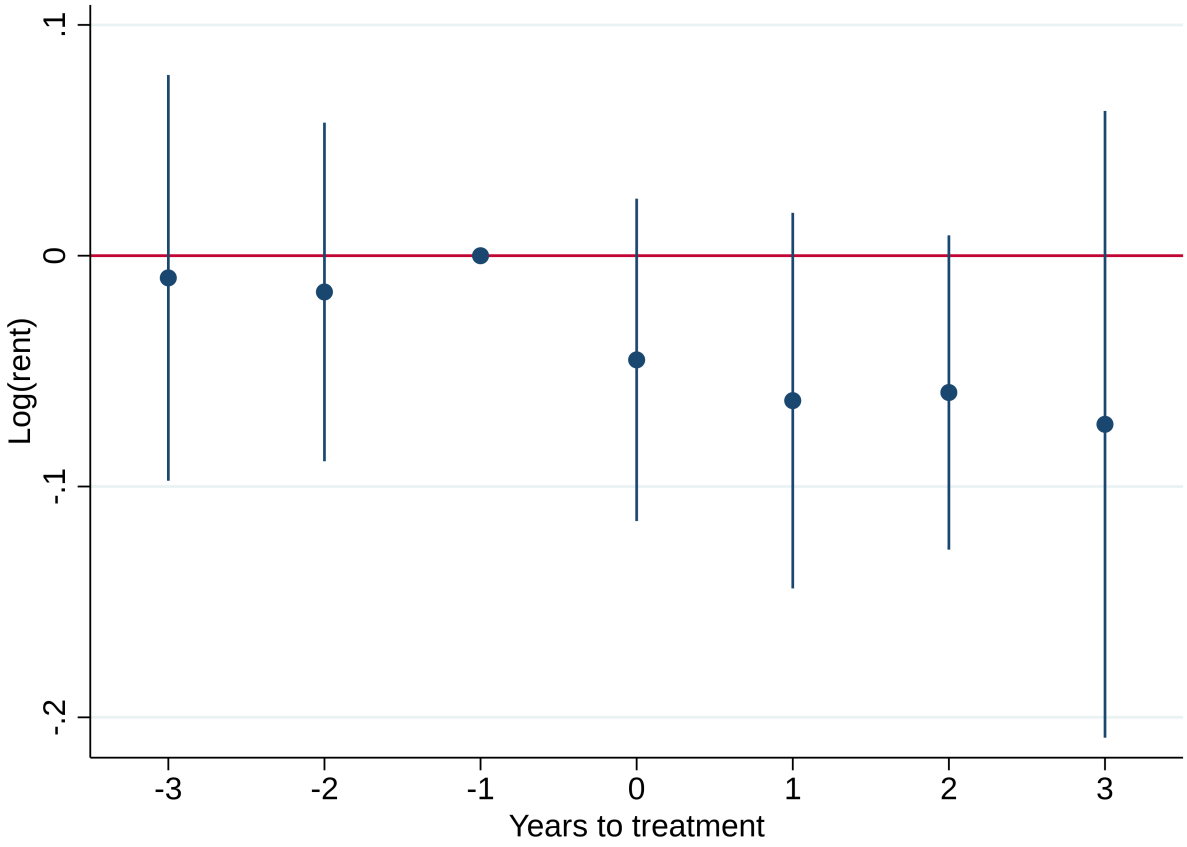

Note: This figure shows the near-far event study of the effect of new buildings on nearby rents. The treatment group is listings within 250m of a new building, and the control group is listings between 250m and 600m of the same buildings. Listings are provided by Zillow™ for years 2013–2018, and we include only the 2015–2016 buildings from our final analysis sample. The specification, shown in Equation 1, includes nearest-building × year and nearest-building × treated fixed effects, as well as controls for bedroom and bathroom counts. Rents are winsorized at the 1st and 99th percentiles. Nearest building-years are weighted equally, and standard errors are clustered at the level of the nearest new building.

Figure 4: Near Versus Near Event Study for Rent Outcome

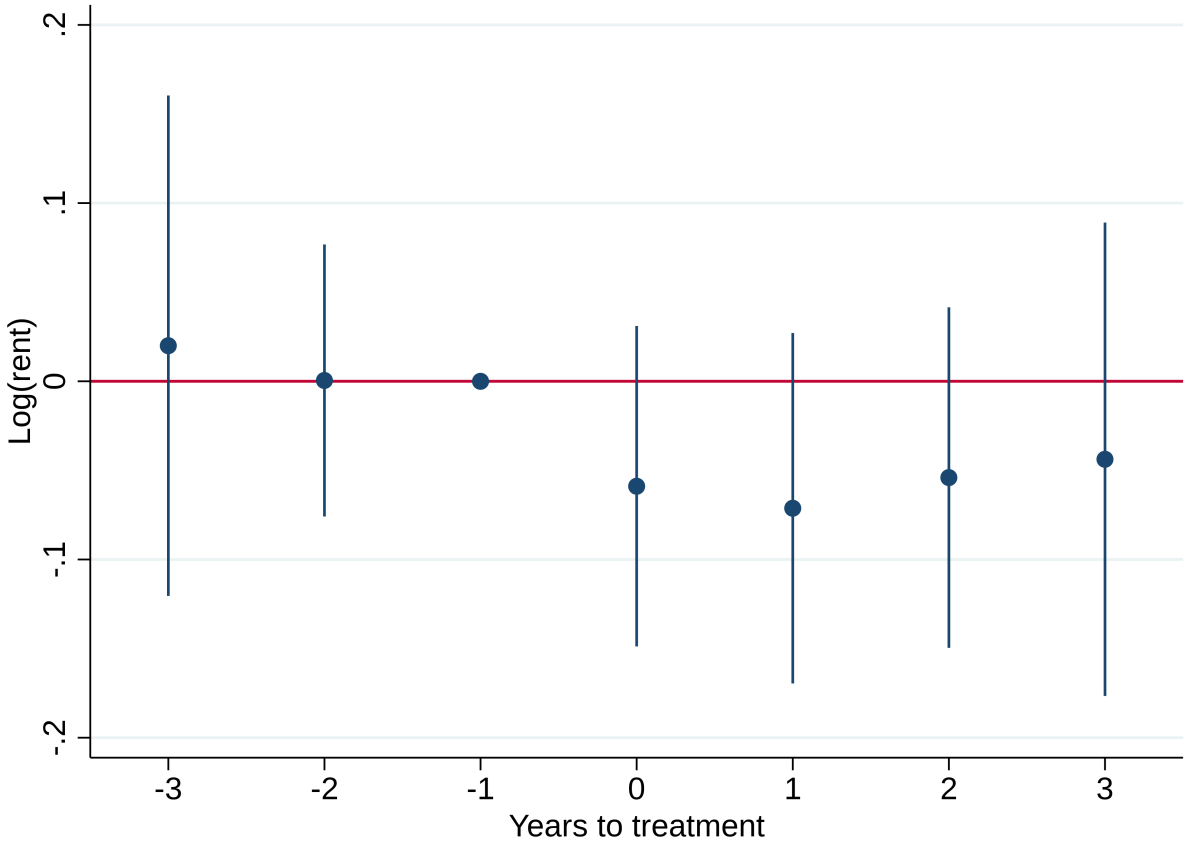

Note: This figure shows the near-near event study of the effect of new buildings on nearby rents. The treatment group is listings within 250m of a building completed in 2015–2016, and the control group is listings within 250m of buildings completed in 2019 (after the sample period). Listings are provided by Zillow™ for years 2013–2018, and we include only the 2015, 2016, and 2019 buildings from our final analysis sample. The specification, shown in Equation 3, include nearest-building and CBSA-year fixed effects, as well as controls for bedroom and bathroom counts. Rents are winsorized at the 1st and 99th percentiles. Nearest building-years are weighted equally, and standard errors are clustered at the level of the nearest new building.

Figure 5: Near Versus Near Event Study for Migration Outcomes

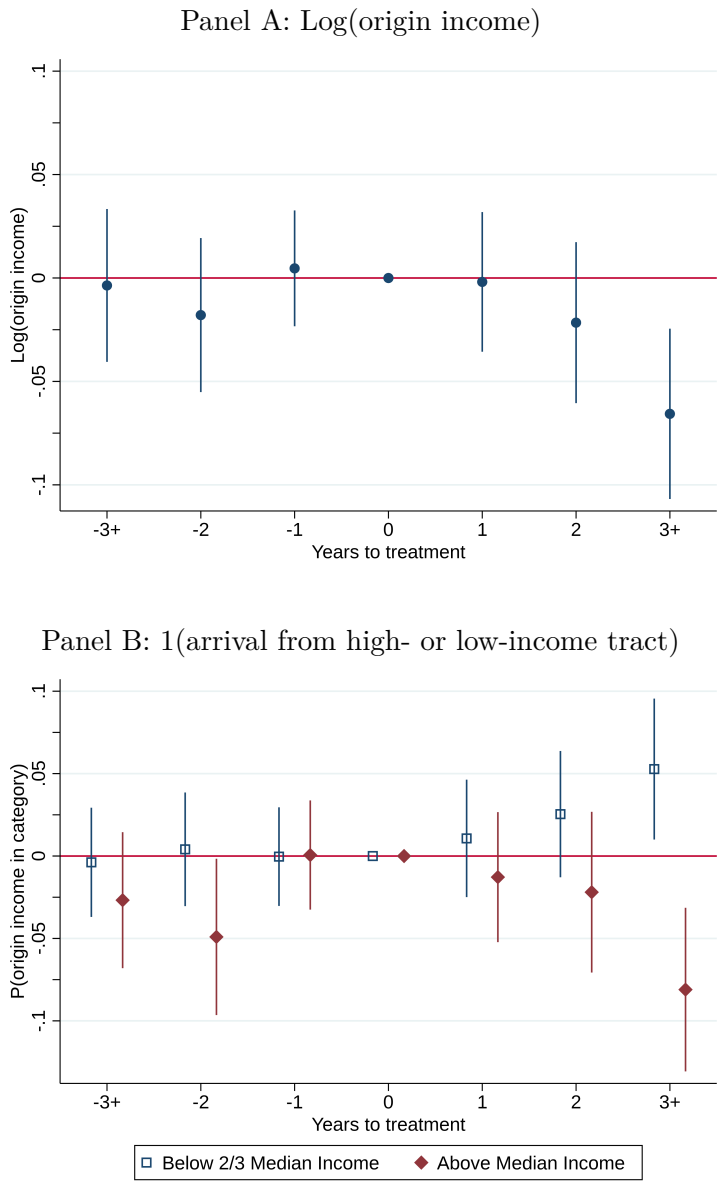

Note: This figure shows the near-near event study of the effect of new buildings on nearby in-migration. The treatment group is arrivals within 250m of a building completed in 2014–2015, and the control group is arrivals within 250m of buildings completed in 2018 (after the sample period). Migration is observed in the Infutor data for years 2011–2017, and we include only the 2014, 2015, and 2018 buildings from our final analysis sample. The specification includes nearest-building and CBSA-year fixed effects Nearest building-years are weighted equally, and standard errors are clustered at the level of the nearest new building.

Figure 6: Near Versus Far Event Study for Migration Outcomes

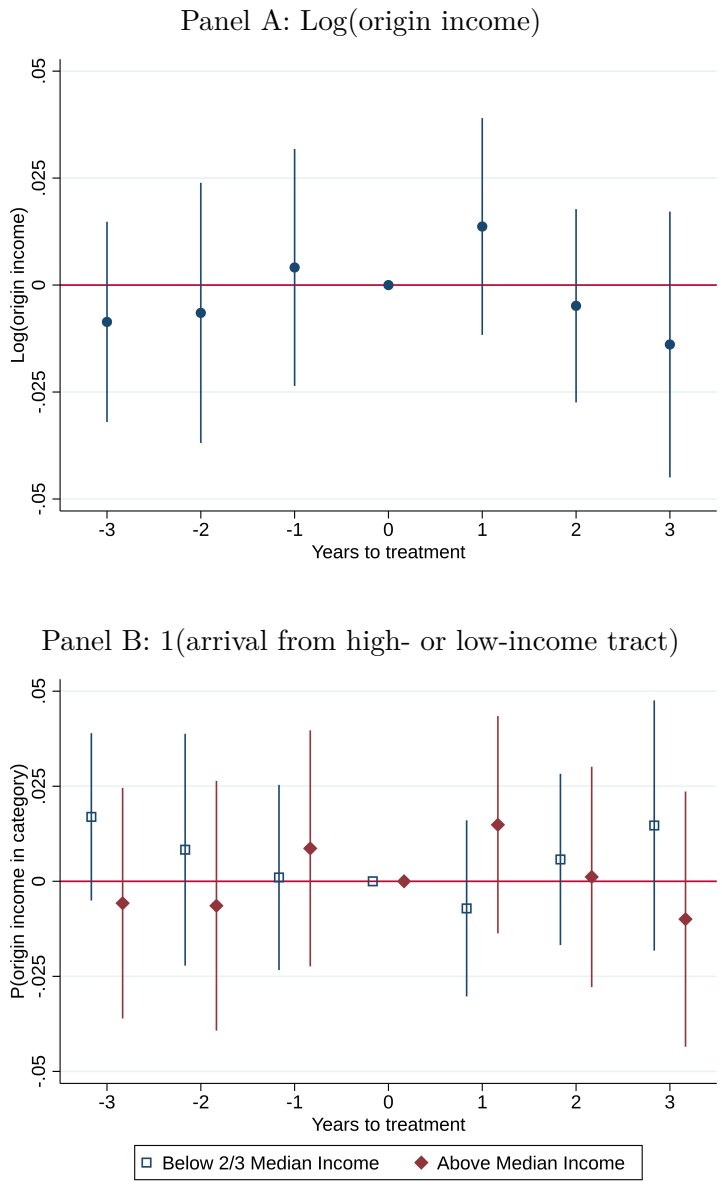

Note: This figure shows the near-far event study of the effect of new buildings on nearby in-migration. The treatment group is arrivals within 250m of a building, and the control group is arrivals between 250m and 600m of the same buildings. Migration is observed in the Infutor data for years 2011–2017, and we include only the 2014 and 2015 buildings from our final analysis sample. The specification includes nearest-building × year and nearest-building × treated fixed effects. Nearest building-years are weighted equally, and standard errors are clustered at the level of the nearest new building.

Table 1: Building Characteristics and Distribution across Cities

	All Incomes 2010-2019	Low-Income 2010-2019	Low-Income 2014-2016	Final Sample
Building Units				
Units (mean)	198	176	168	166
Units (25th pctile)	85	79	76	68
Units (50th pctile)	162	139	131	121
Units (75th pctile)	287	249	250	256
Nearby Rents				
Rent 1 br (mean)	1,784	1,609	1,568	1,539
Rent 1 br (25th pctile)	1,356	1,290	1,284	1,262
Rent 1 br (50th pctile)	1,633	1,528	1,508	1,514
Rent 1 br (75th pctile)	2,045	1,801	1,739	1,740
Rent (mean)	2,129	1,917	1,875	1,855
Rent (25th pctile)	1,667	1,572	1,597	1,601
Rent (50th pctile)	1,942	1,799	1,784	1,807
Rent (75th pctile)	2,441	2,098	2,037	2,052
Buildings by City (column percentages)				
Atlanta	7	4	5	4
Austin	10	12	14	13
Brooklyn	13	13	11	13
Chicago	8	3	2	4
Denver	9	10	10	15
Los Angeles	11	13	11	15
Philadelphia	3	3	2	1
Portland	9	11	11	10
San Francisco	4	3	2	1
Seattle	17	20	25	20
Washington, D.C.	7	7	6	3
Observations	1,483	823	253	92

Note: Distributions of units in buildings, rents in Zillow™ listings near buildings (within 600 meters), and city locations of buildings for different samples of buildings. Samples of buildings are described in detail in Section 1.1 of the main text. Building locations and unit counts are provided by Real Capital Analytics.

Table 2: Building Neighborhood Characteristics

	All Incomes		Low-Income			
	No Building	Some Building	No Building	Some Building	2014-6 Building	Final Sample
Household Income						
2000 ($)	57,879	55,240	46,575	43,376	44,278	42,732
2010 ($)	58,328	61,513	45,116	45,201	45,722	44,949
2017 ($)	60,330	76,332	44,724	53,369	53,631	52,544
2000 to 2010 (pct)	-1	11	-4	5	4	5
2010 to 2017 (pct)	2	22	-1	17	16	16
College Degree						
2000 (pct)	25	39	17	30	32	27
2010 (pct)	30	50	22	40	42	37
2017 (pct)	34	60	26	49	51	48
2000 to 2010 (pp)	5	12	4	10	10	10
2010 to 2017 (pp)	4	9	4	9	9	10
Rent						
2000 ($)	1,001	992	884	872	884	845
2010 ($)	1,168	1,199	1,045	1,002	995	954
2017 ($)	1,271	1,506	1,111	1,238	1,227	1,175
2000 to 2010 (pct)	16	19	17	14	12	12
2010 to 2017 (pct)	7	22	6	21	21	20
Observations	3,218	1,483	2,375	823	253	92

Note: Means of the characteristics of the neighborhoods (census tracts) where our new buildings were built. The left two columns include all central city tracts in our sample cities, while the right four only include those with median household income below the CBSA median. “No Building” denotes neighborhoods that did not receive any of our sample buildings between 2010 and 2019. “Some building" includes neighborhoods with at least one building, “2014-6 Building" is neighborhoods with a building completed between 2014 and 2016, and “Final Sample" is tracts that contain a building in the final analysis sample described in Section 1.1. For columns with some building, the column means are weighted by the number of buildings in each neighborhood. Information on building locations is provided by Real Capital Analytics, and demographic information comes from the Census 2000 Long Form (“2000”), American Community Survey 2006-2010 5-Year Estimates (“2010”), and American Community Survey 2013-2017 5-Year Estimates (“2017”).

Table 3: Treatment versus Control Buildings in Rent Near-Near Analysis

Dependent Variable	β	*t-stat*
Median household income	6025	1.3
2BR Rent	169	1.6
Pct. White	0.06	1
Pct. Black	-0.029	0.6
Pct. College	0.09	1.8
Pct. Poverty	-0.05	1.6
Pct. 200k+	0.005	0.5
Pct. Owner-Occupied	-0.03	-0.6
Pct. Vacant	0.012	1
Population	-184	0.3

Note: This table shows differences between the neighborhoods around the treatment (2015–2016 completion) and control (2019) buildings in the near-near rental analysis. Each row shows results from a regression of the dependent variable shown in the first column on a treatment indicator and a set of CBSA fixed effects. β is the coefficient on the treatment indicator, and t is the t-statistic on that coefficient

Table 4: Baseline Difference-in-Differences Results for Rent Outcomes

	Near versus far	Near versus near	Triple-difference
After*within 250	-0.049		
(S.E.)	(0.021)		
After*treated building		-0.062	
(S.E.)		(0.037)	
After*within 250*treated building			-0.071
(S.E.)			(0.033)
Treated buildings	46	44	42
Control buildings		20	19
Listing observations	56,000	20,400	56,800

Note: This table shows baseline difference-in-differences results for rents. The first column shows the near-far specification shown in Equation 2, where the treatment group is listings within 250m of a building completed in 2015–2016, and the control group is listings between 250m and 600m of the same buildings. The second column shows the near-near specification, shown in Equation 4, in which the treatment group is listings within 250m of a building completed in 2015–2016, and the control group is listings within 250m of buildings completed in 2019 (after the sample period). The third column shows the triple-difference specification from Equation 5, which compares the near-far gap in the 2015–2016 and 2019 buildings. Listings are provided by Zillow™ for years 2013–2018, and we include only new buildings that meet the restrictions of our final analysis sample. All specifications include controls for bedroom and bathroom counts, and rents are windsorized at the 1st and 99th percentiles. Building-years are weighted equally, and standard errors are clustered at the level of the nearest new building.

Table 5: Baseline Difference-in-Differences Results for Migration Outcomes

	Near versus far	Near versus near	Triple-difference
Panel A: Log(origin income)			
After*within 250	-0.004		
(S.E.)	(0.01)		
After*treated building		-0.03	
(S.E.)		(0.017)	
After*within 250*treated building			-0.027
(S.E.)			(0.014)
Panel B: 1(above median origin)			
After*within 250	-0.0003		
(S.E.)	(0.01)		
After*treated building		-0.023	
(S.E.)		(0.021)	
After*within 250*treated building			-0.019
(S.E.)			(0.017)
Panel C: 1(<2/3 median origin)			
After*within 250	0.002		
(S.E.)	(0.008)		
After*treated building		0.034	
(S.E.)		(0.017)	
After*within 250*treated building			0.027
(S.E.)			(0.013)
Treated buildings	66	66	66
Control buildings		72	72
Arrival observations	119,400	57,400	186,700

Note: This table shows baseline difference-in-differences results for migration. Each panel uses the dependent variable indicated in the heading. The first column shows the near-far specification, where the treatment group is arrivals within 250m of a building completed in 2014-2015, and the control group is arrivals between 250m and 600m of the same buildings. The second column shows the near-near specification, in which the treatment group is arrivals within 250m of a building completed in 2014–2015, and the control group is arrivals within 250m of buildings completed in 2018 (after the sample period). The third column shows the triple-difference specification, which compares the near-far gap in the 2014–2015 and 2018 buildings. Migration data for 2011–2017 is from Infutor, and we include only new buildings that meet the restrictions of our final analysis sample. Building-years are weighted equally, and standard errors are clustered at the level of the nearest new building.

A Appendix

A.1 Empirical Derivative Estimation Details

In this section, we describe our approach to estimate a continuous treatment effect of new buildings, which follows Diamond and McQuade (2019). While full details on the methodology are available in that paper, we provide a summary here. The basic idea is to calculate how rents changed with distance from the development site, holding constant rental listing timing and proximity as closely as possible. The equation of interest is:

$$log(rent_{jt}) = \tilde{m}_{\mathbf{Y}}(r_j, \tau_j) + \phi_l(r_j, \theta_j) + \psi_l(\theta_j, t) + \epsilon_{jt} \tag{6}$$

where $rent_{jt}$ is the listed rent of apartment unit j in year t, unit locations are defined in polar coordinates (r_j, θ_j), and τ represents years to building completion. $\tilde{m}_{\mathbf{Y}}(r_j, \tau_j)$ is the semiparametric function of interest, which describes the relationship between rents and distance at different years to the building's completion. Our main results for the exercise are graphs of this function. $\phi_l(r_j, \theta_j)$ represent "neighborhood-specific" fixed effects, and $\psi_l(\theta_j, t)$ are distinct time trends for neighborhood l (which could vary by θ).

To estimate the equation, we calculate the change in prices over r by differencing over pairwise transactions that vary in r but are as close in t and θ as possible. We then average over these pairwise calculations, and smooth the estimates with a Nadaraya-Watson kernel estimator. This yields the estimated derivative of the $\tilde{m}_{\mathbf{Y}}(r_j, \tau_j)$ function that we plot in our results.

We depart from Diamond and McQuade chiefly by first demeaning rents according to bedroom and bathroom counts, which is necessary because of the large variation in rental units within a short distance. In addition, since we study a shorter time span than Diamond and McQuade and use a smaller outer radius, we choose a smaller Nadaraya-Watson bandwidth. We set $h_{t,n} = 3$ instead of $h_{t,n} = 5$ as our time bandwidth, and $h_{r,n} = 0.2$ instead of $h_{r,n} = 0.3$ as our distance bandwidth. Additionally, since listings are relatively more common than housing price transactions, we use 12 pairwise transactions for each calculation of the derivative.

A.2 Additional Figures and Tables

Figure A.1: 2018 Rental Listings Near Chicago Buildings

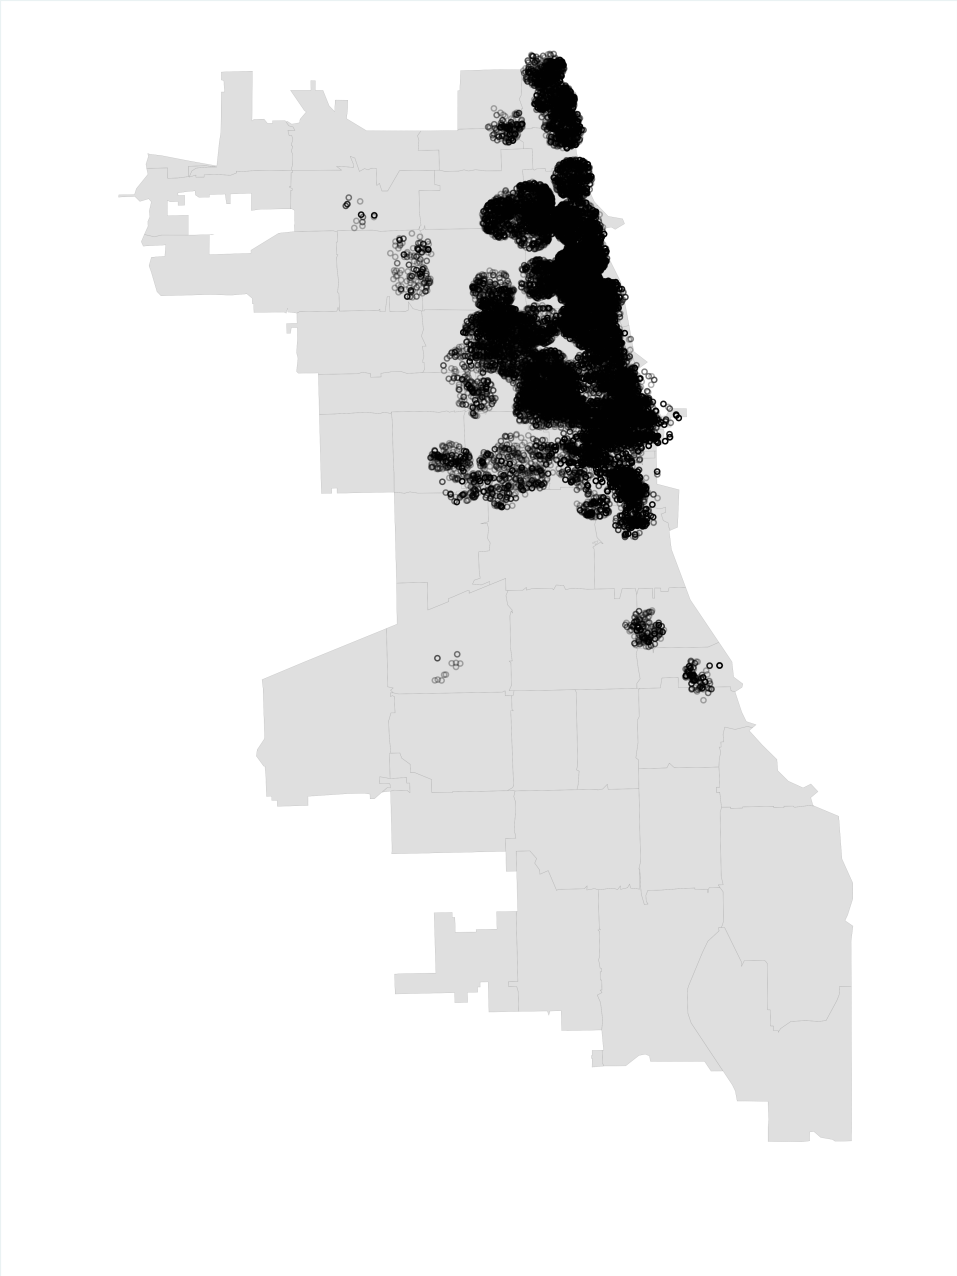

Note: This figure shows the location of 2018 Zillow™ listings in Chicago that are within 800 meters of a sample building completed between 2010 and 2019.

Figure A.2: Rent vs. Distance to Nearest Building

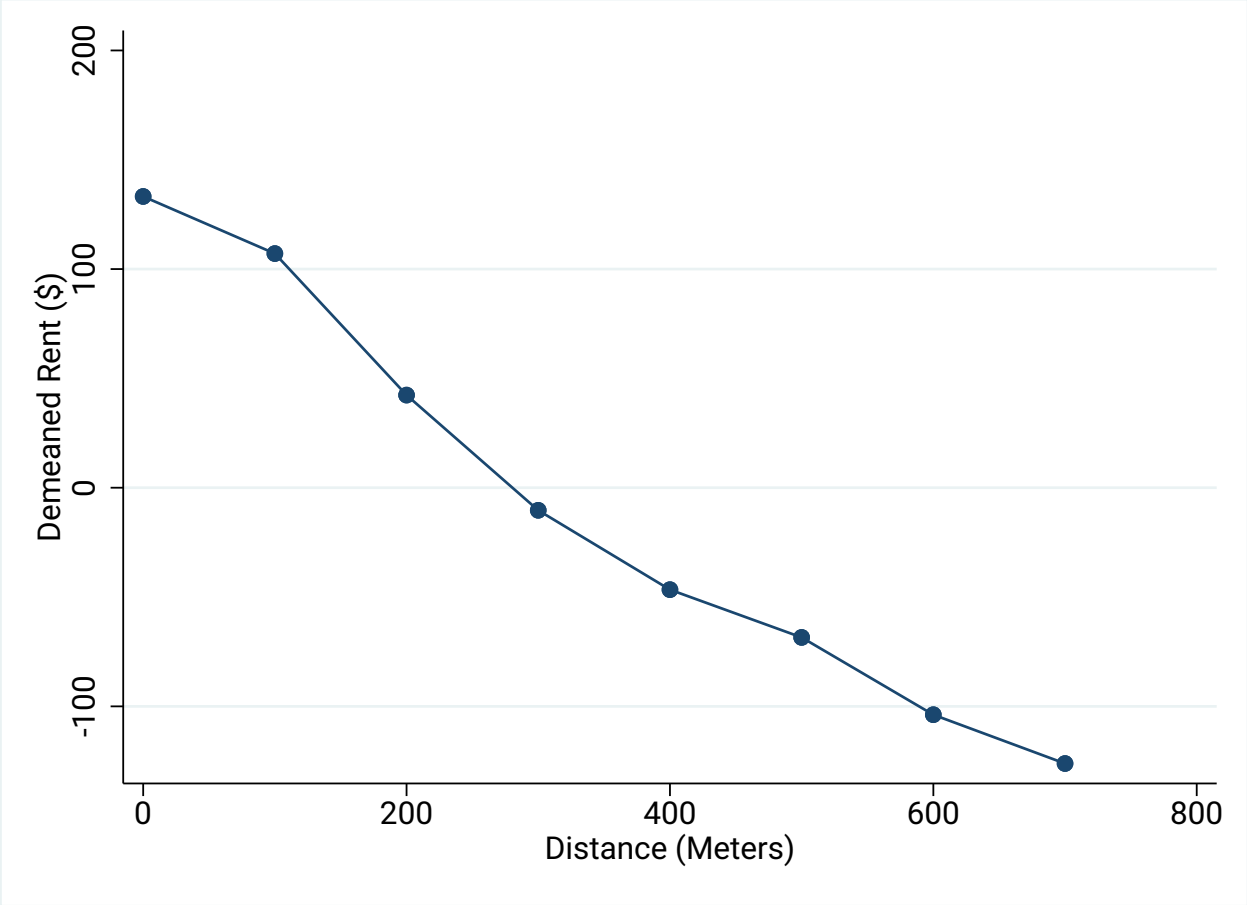

Note: This figure shows rent versus distance to the nearest new building for units in the Zillow™ sample. Rent is de-meaned at the CBSA-year level.

Figure A.3: Infutor vs. Census Population (Census Tract Level)

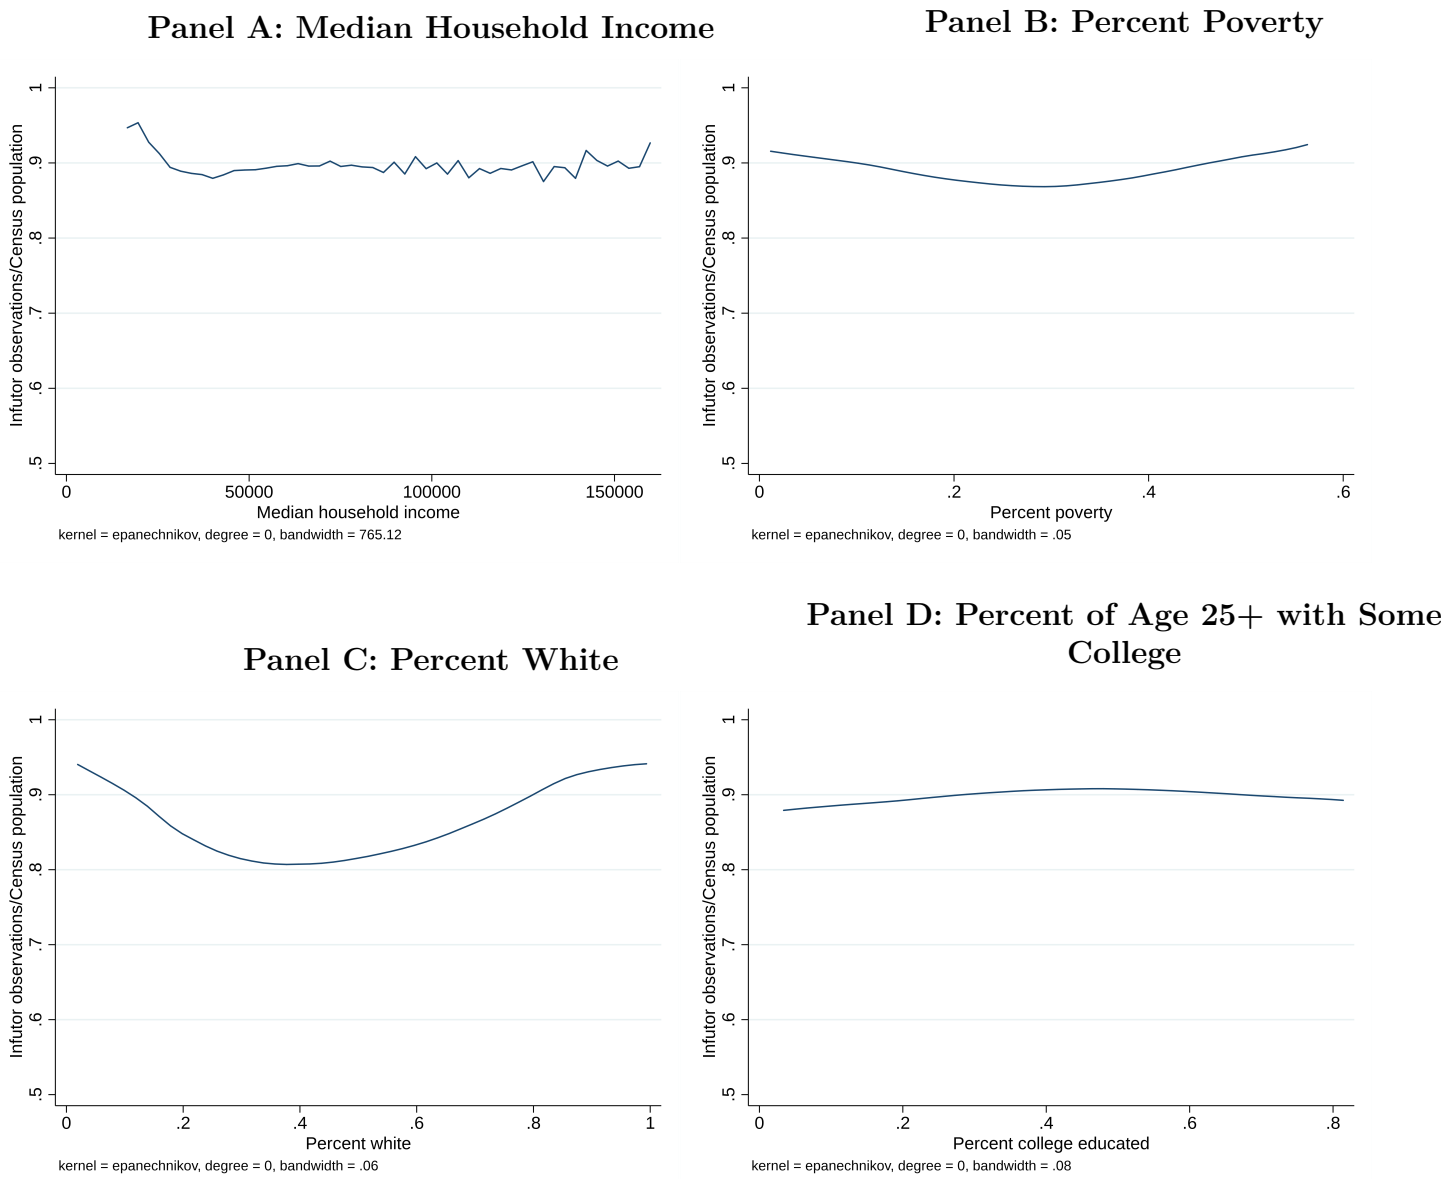

Note: Each panel plots a local polynomial regression of Infutor coverage (measured as the ratio of Infutor observations to census over-25 population) in a census tract versus the tract characteristic in the heading. Census figures are drawn from the 2013–2017 ACS.

Figure A.4: Infutor vs. Census Migration Rates (County Level)

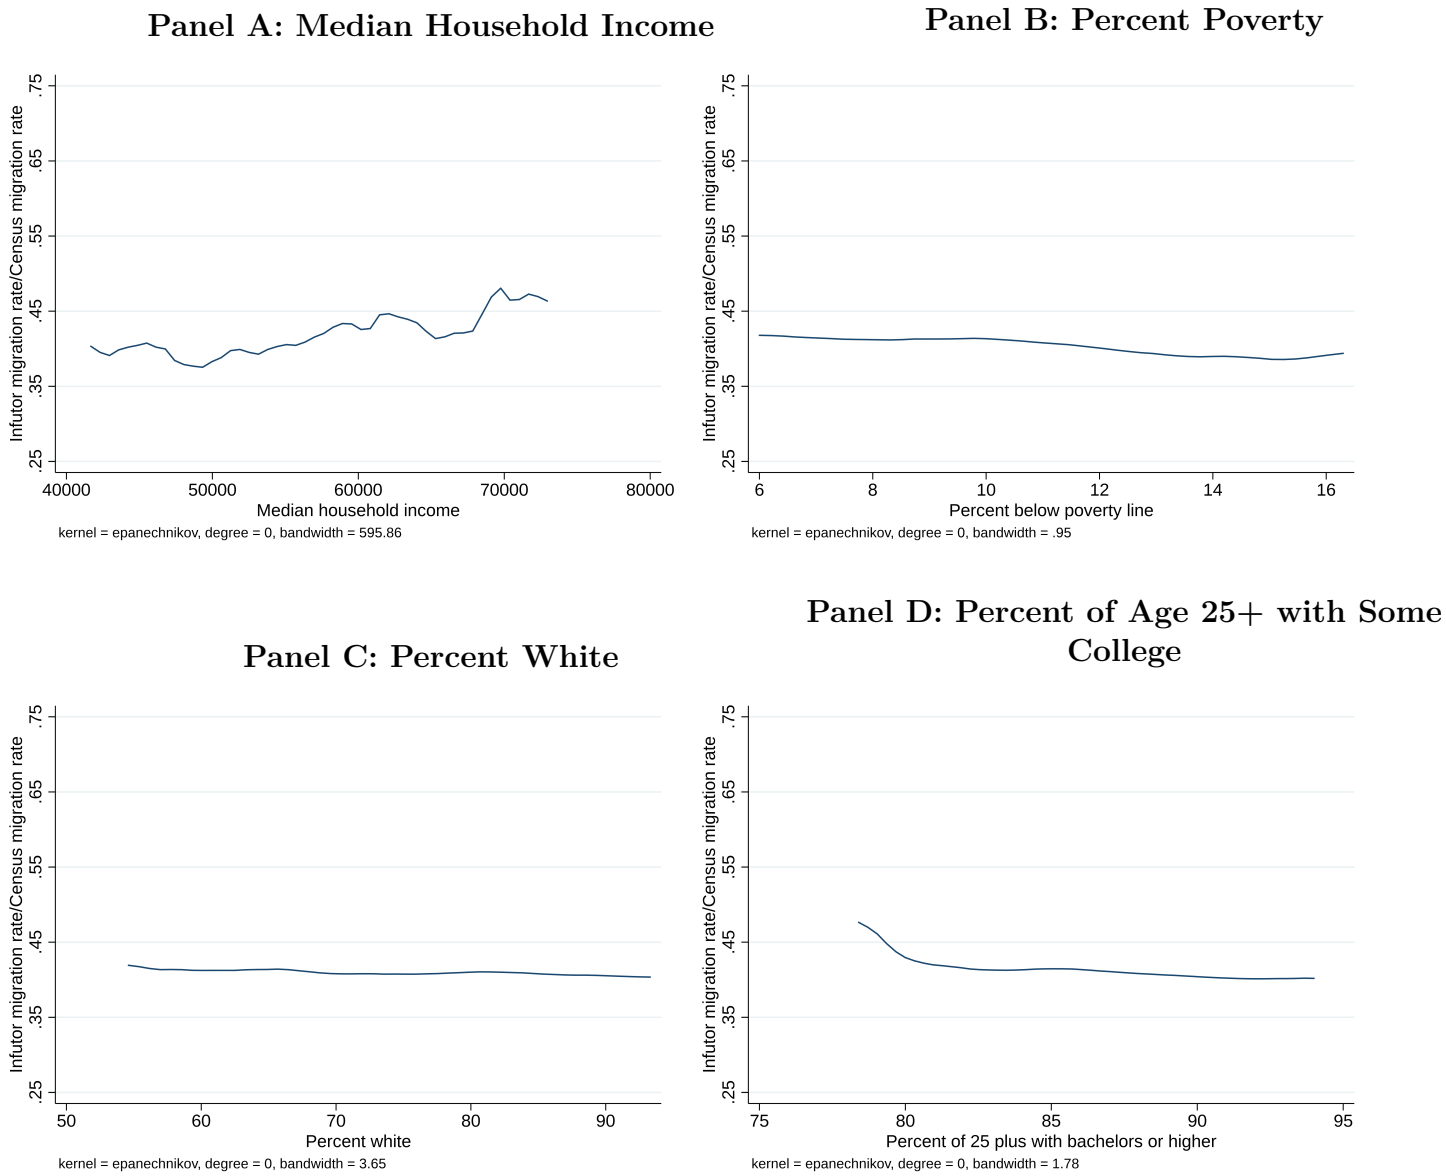

Note: Each panel plots a local polynomial regression of the ratio of Infutor to census move rates (measured at the county level) against county characteristics. County characteristics and move rates are drawn from the 2013–2017 ACS.

Figure A.5: Example Near-Far Samples

Panel A: Portland

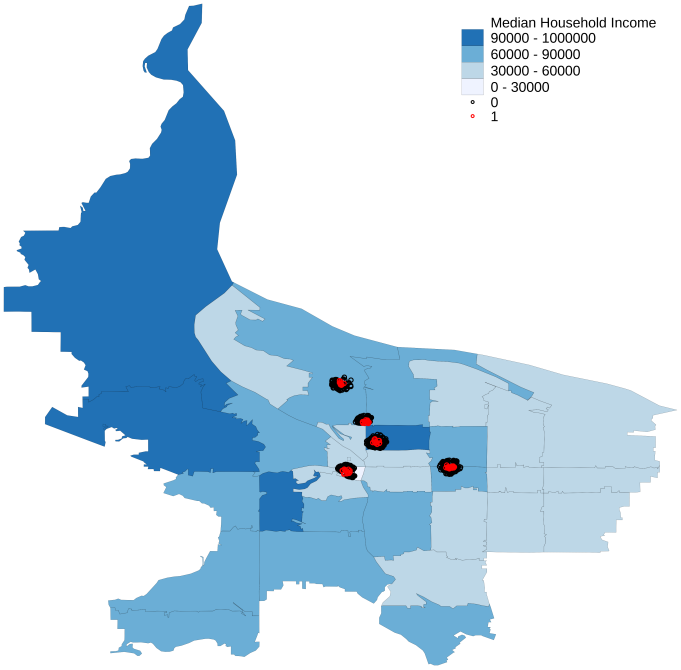

Panel B: Brooklyn

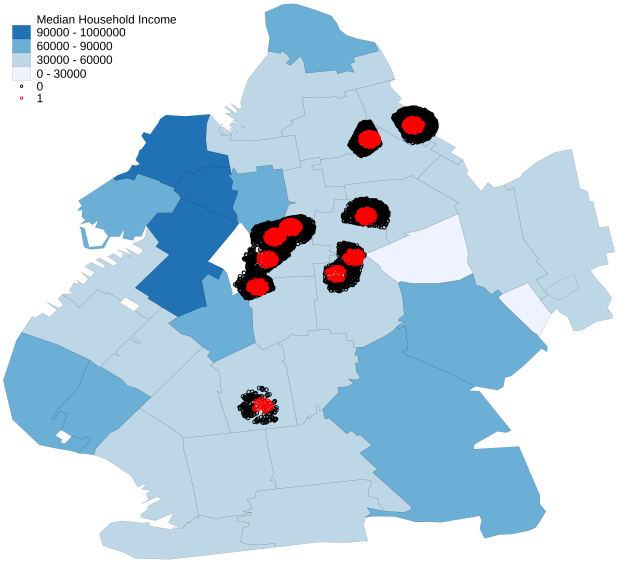

Note: This figure shows the near-far rental specification treatment (within 250 meters) and control groups (250-600 meters) in Brooklyn and Portland. Only 2015–2016 buildings that meet the sample criteria in Section 1.1 are included, and listings are associated with their nearest such building. Polygons represent zip codes and are colored according to median household income in the 2013–2017 ACS.

Figure A.6: Near-Far Event Study for Rent Outcome (400m Treatment Radius)

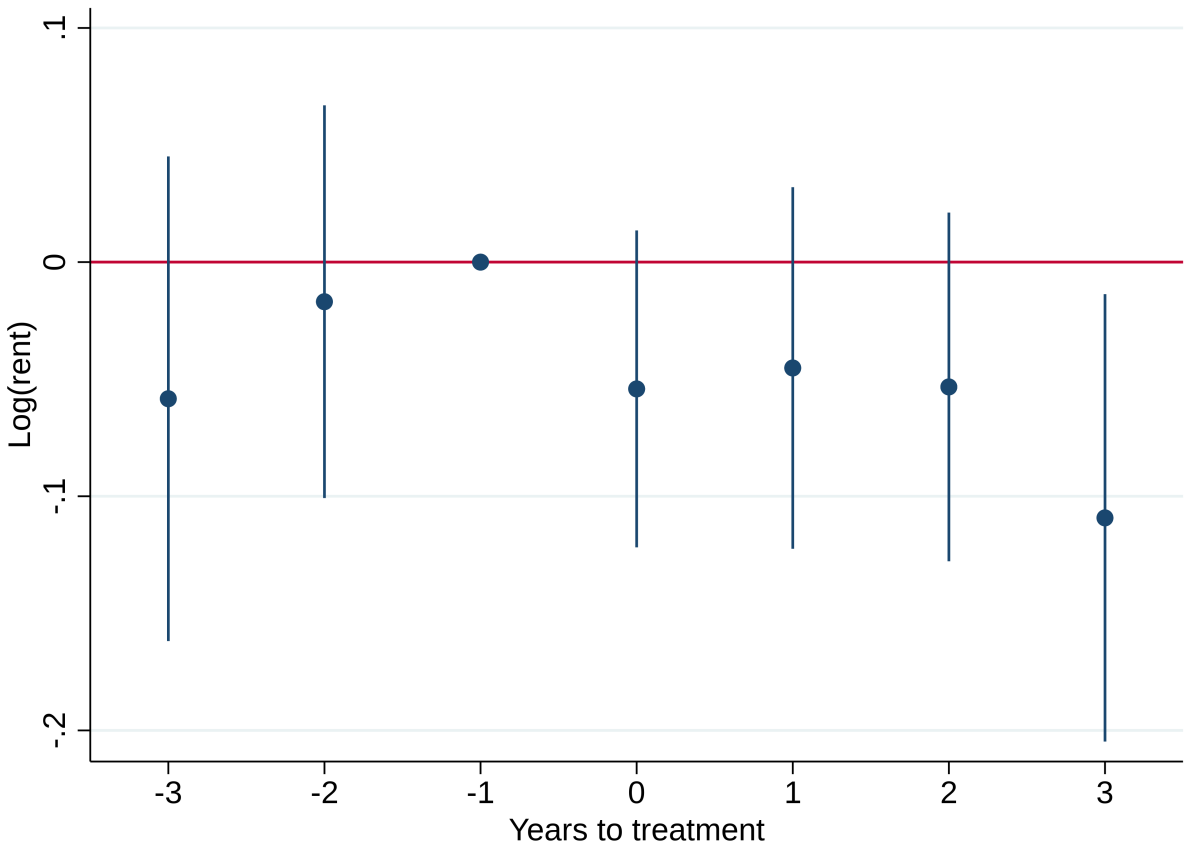

Note: This figure shows the near-far event study of the effect of new buildings on nearby rents. The treatment group is listings within 400 meters of a building completed in 2015–2016, and the control group is listings between 400m and 600m of the same buildings. In addition, we expand the isolation requirement for new buildings from the baseline to 400 meters to match the larger treatment radius. Otherwise, the specification is as described in Equation 3, and other details are identical to Figure 3.

Figure A.7: Near-Near Event Study for Rent Outcome (400m Treatment Radius)

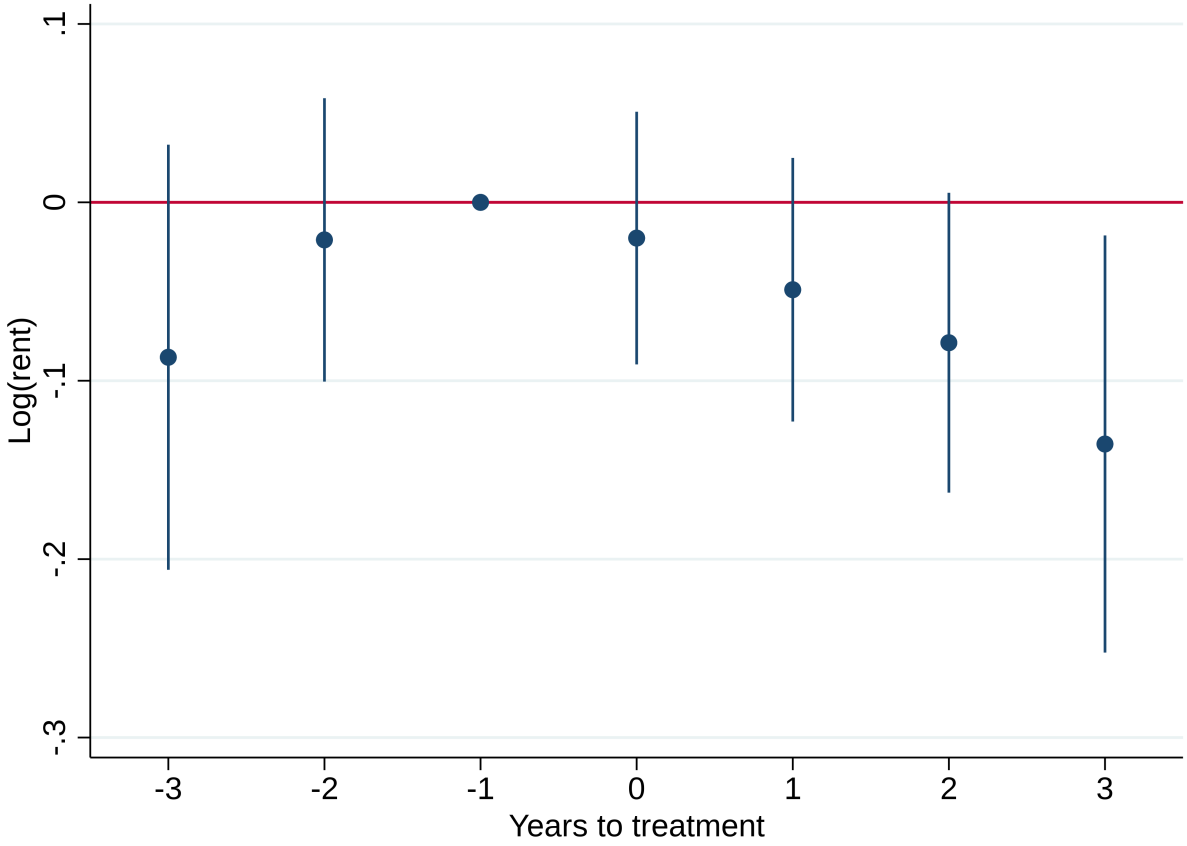

Note: This figure shows the near-near event study of the effect of new buildings on nearby rents. The treatment group is listings within 400m of a building completed in 2015–2016, and the control group is listings within 400m of buildings completed in 2019 (after the sample period). In addition, we expand the isolation requirement for new buildings from the baseline to 400 meters to match the larger treatment radius. Otherwise, the specification is as described in Equation 4, and other details are identical to Figure 4.

Figure A.8: Empirical Derivative Results

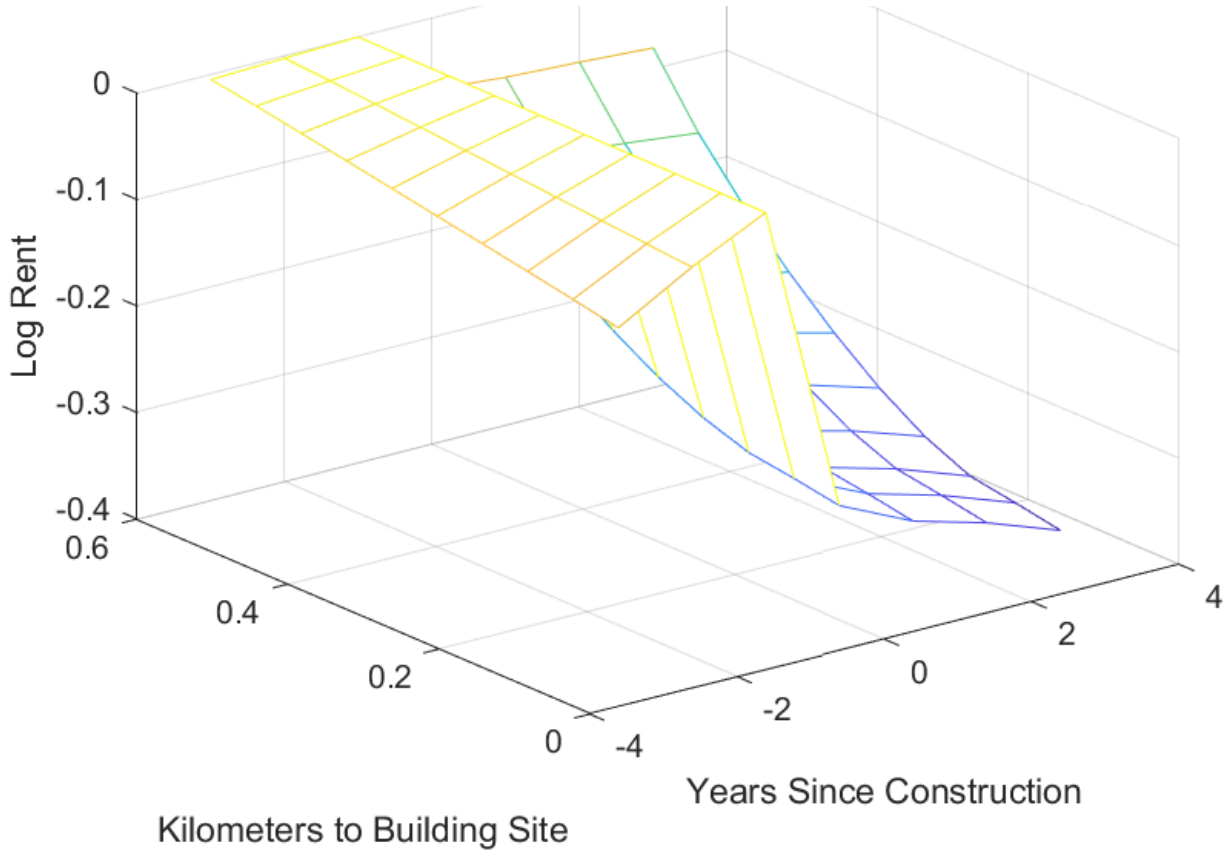

Note: This surface represents the treatment effect of new buildings at every combination of distance from a new building and year since building completion. Further details are provided in the Appendix.

Figure A.9: Near-Far Event Study Robustness (Rent Outcome)

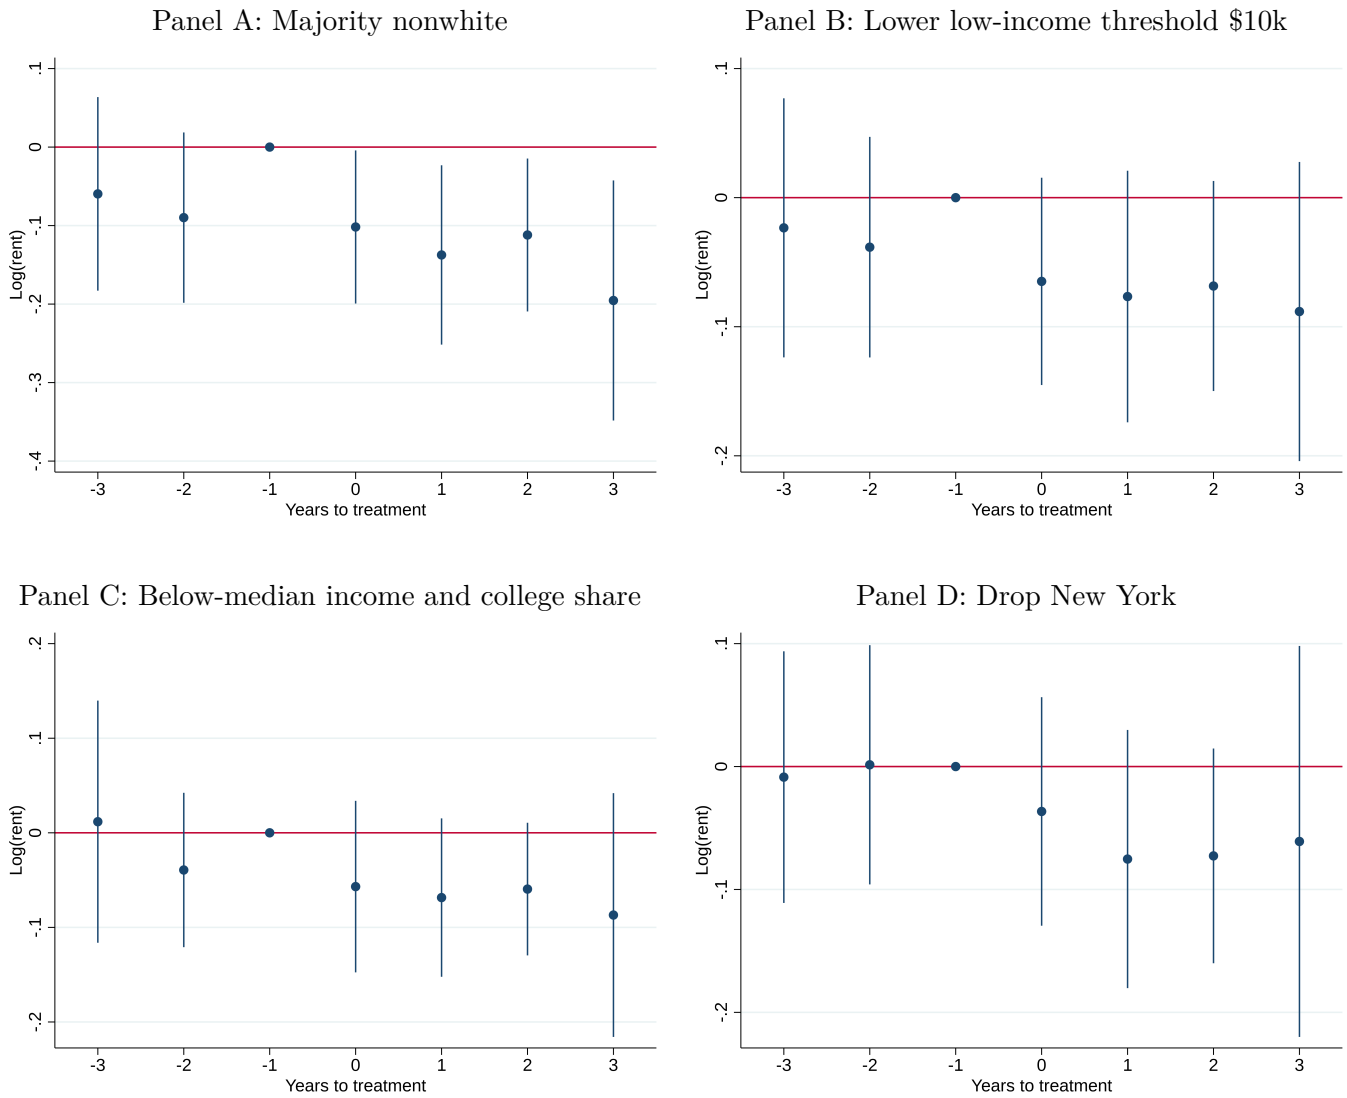

Note: Each panel repeats the baseline near-far event study shown in Figure 3 with a change to the sample. Panel A drops buildings in tracts that are over 50 percent white, and Panel B lowers the income threshold by $10,000. Panel C requires that both tract income and college share be below the CBSA median, while Panel D drops observations in New York City.

Figure A.10: Near-Near Event Study Robustness (Rent Outcome)

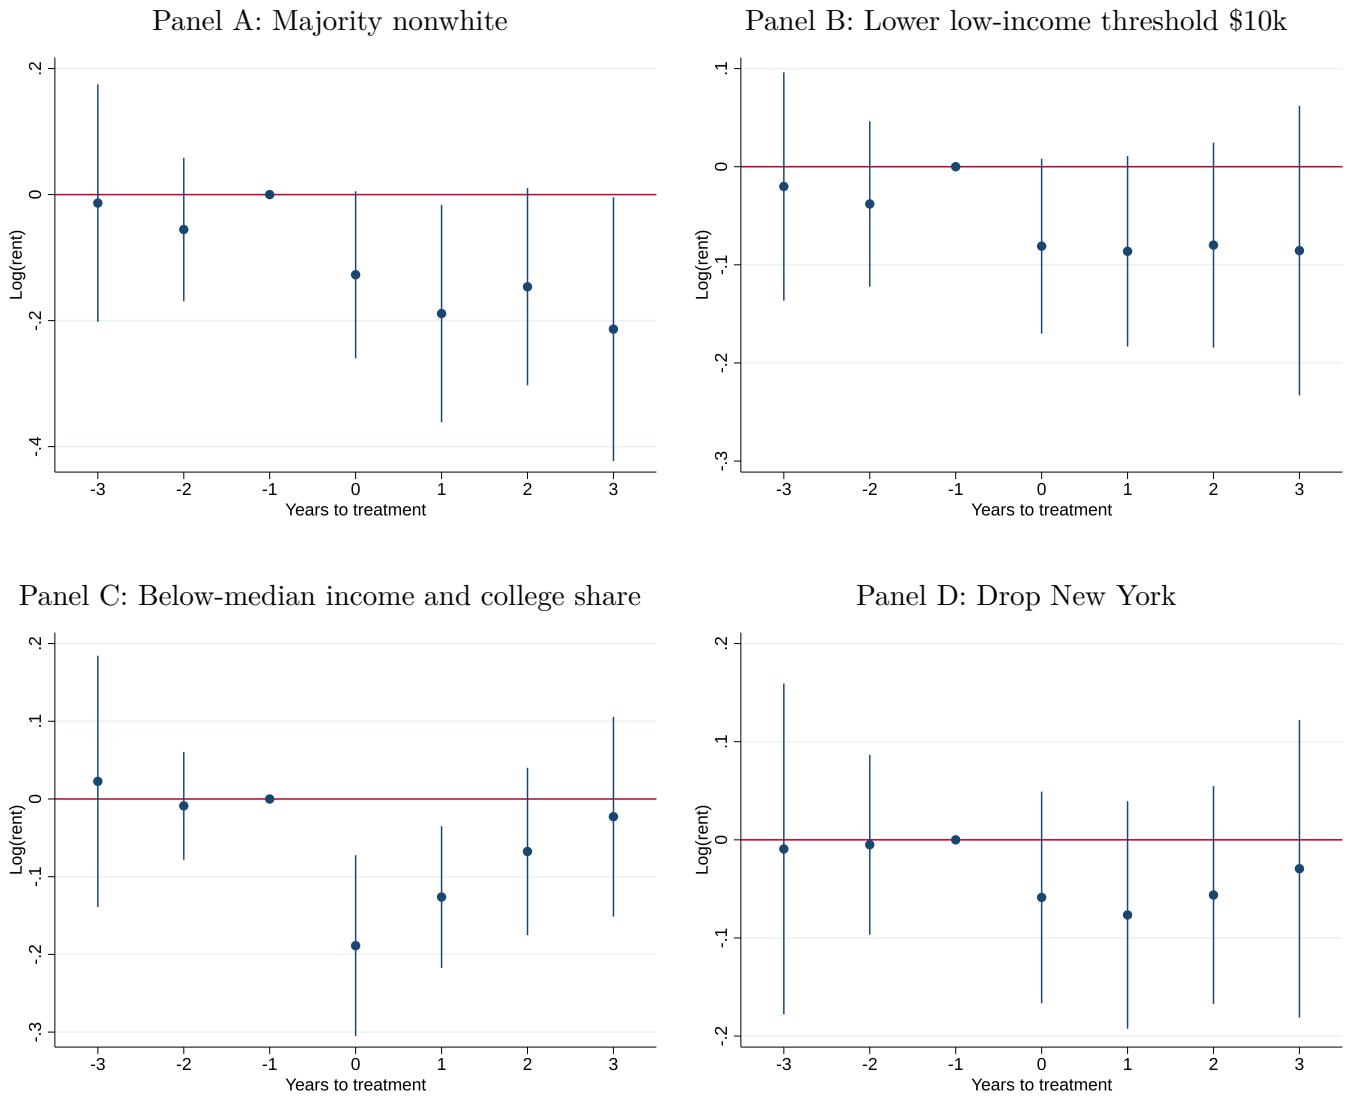

Note: Each panel repeats the baseline near-near event study shown in Figure 4 with a change to the sample. Panel A drops buildings in tracts that are over 50 percent white, and Panel B lowers the income threshold by $10,000. Panel C requires that both tract income and college share be below the CBSA median, while Panel D drops observations in New York City.

Figure A.11: Longer-Run Near-Far Effect for 2010–2013 Buildings

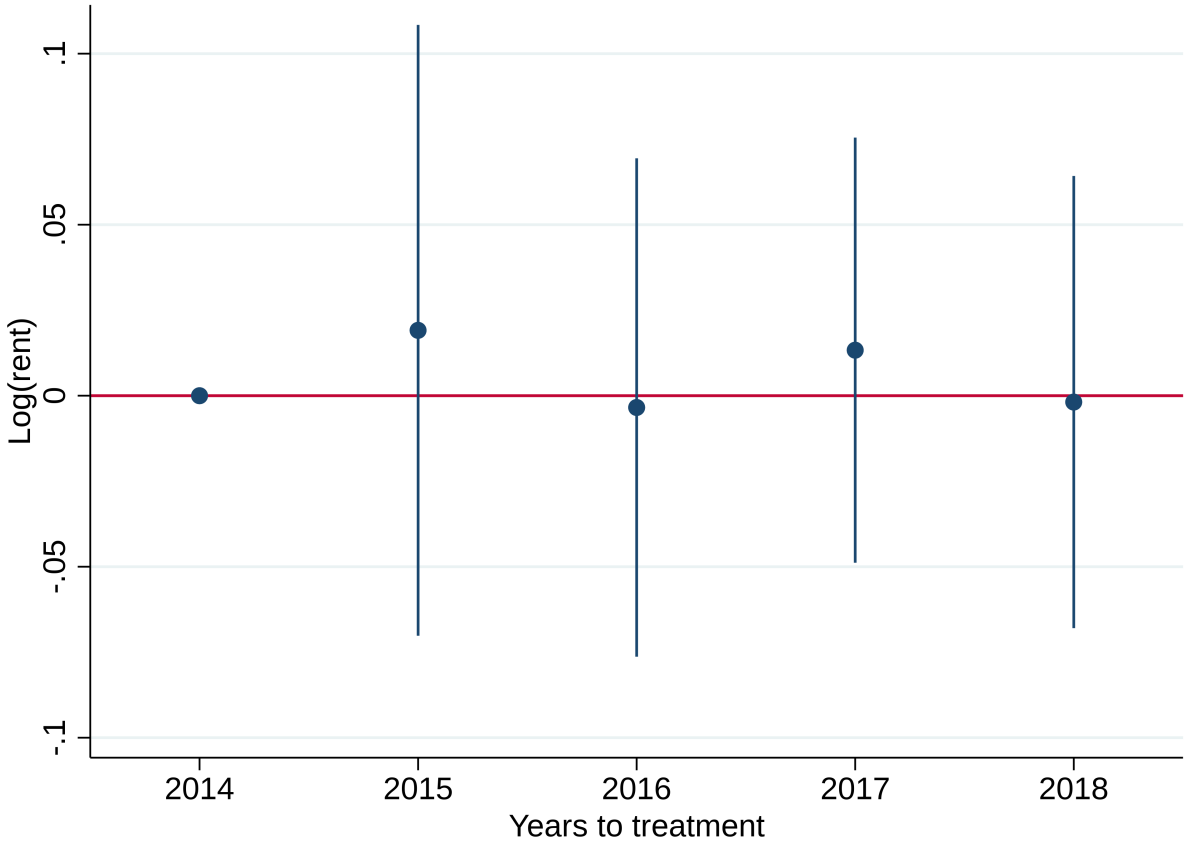

Note: In this figure, we repeat the main near-far event study specification including only buildings completed before 2014. We restrict to the post-period and include year × treatment dummies instead of years-to-treatment dummies so that the sample composition underlying the dummies does not change.

Figure A.12: Longer Near-Far Pretrends for 2019 Buildings

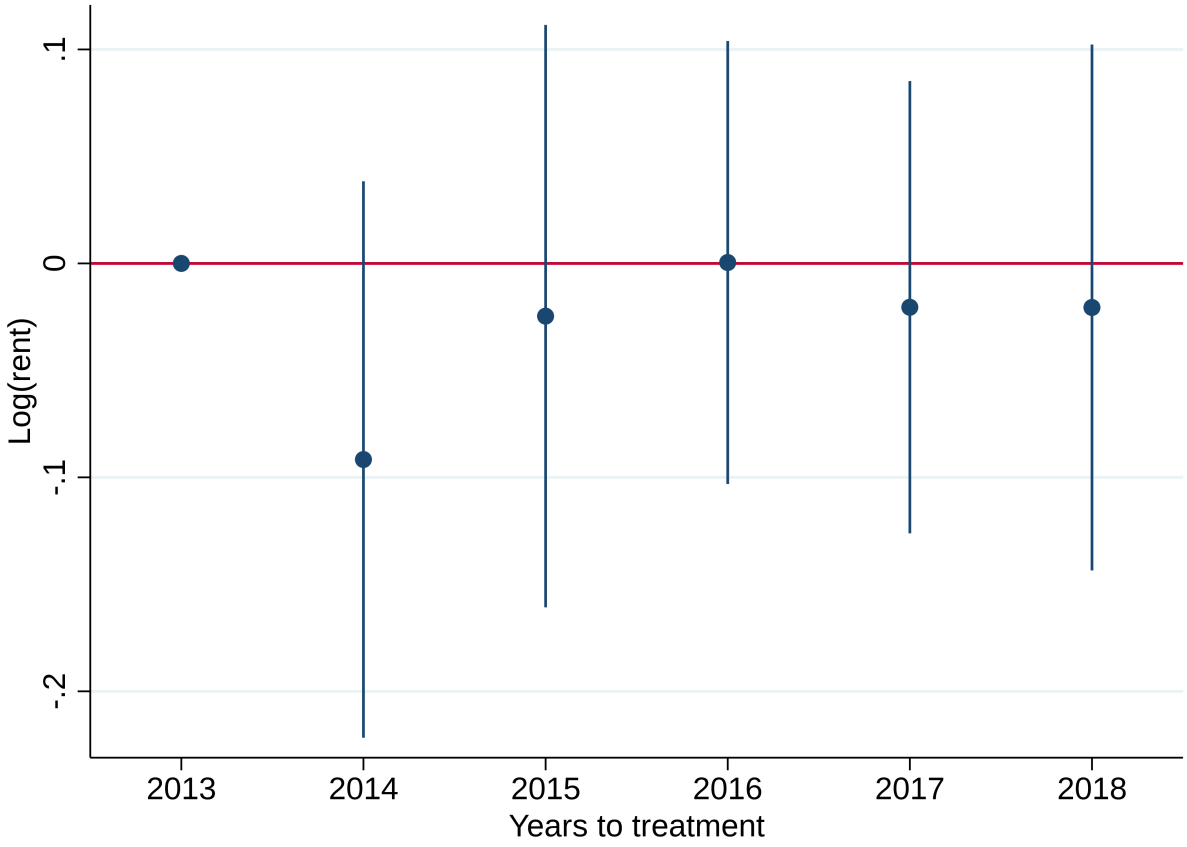

Note: In this figure, we repeat the main near-far event study specification including only buildings completed in 2019. We restrict to the pre-period and note that the year × treatment dummies shown in the figure are equivalent to years-to-treatment dummies because there is only one year of treatment buildings.

Figure A.13: Hyperlocal Price Pretrends in 2018-2019 Building Sites

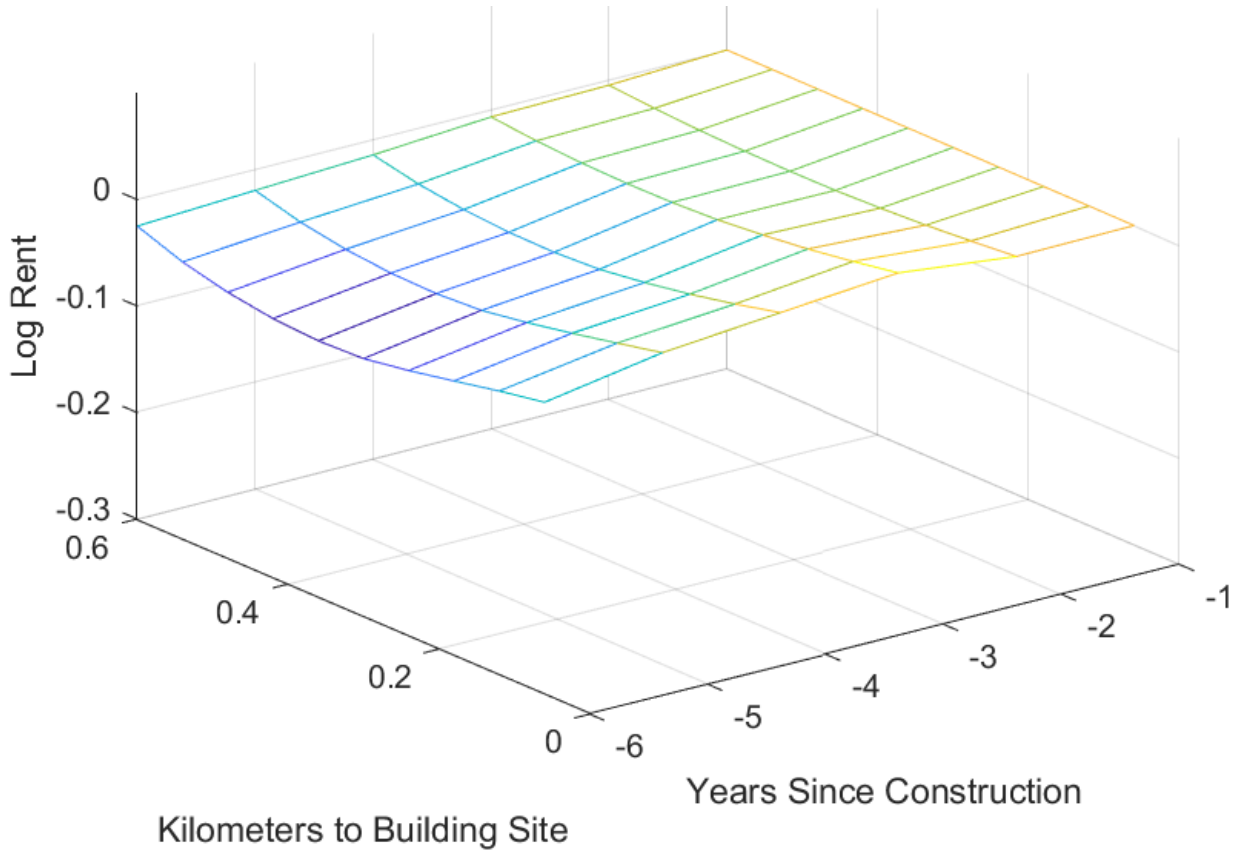

Note: This surface represents the pre-period treatment effect of new buildings completed in 2019. Further details are included in the Appendix.

Figure A.14: Near-Far Event Studies for Rent Outcome (All New Buildings)

Panel A: 250m Treatment Radius

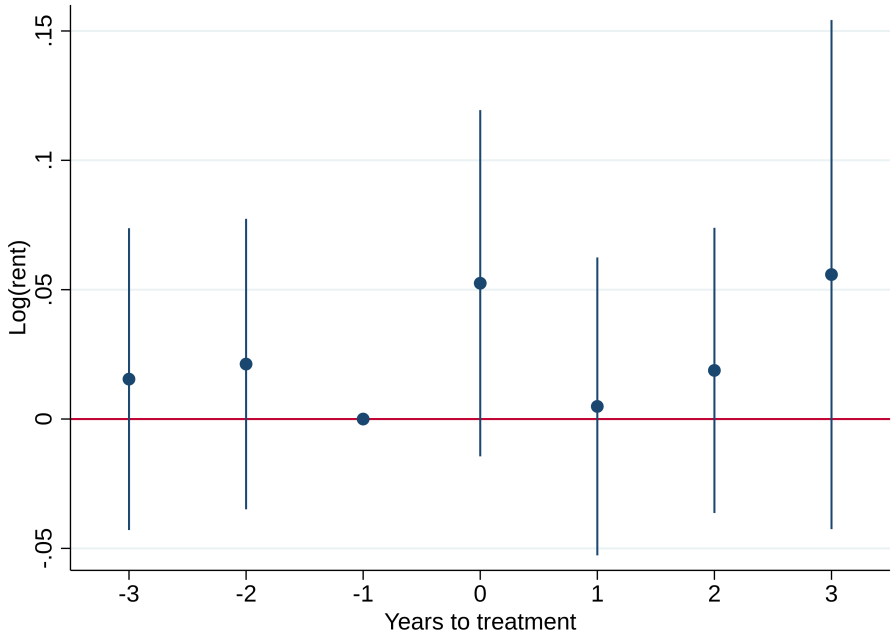

Panel A: 400m Treatment Radius

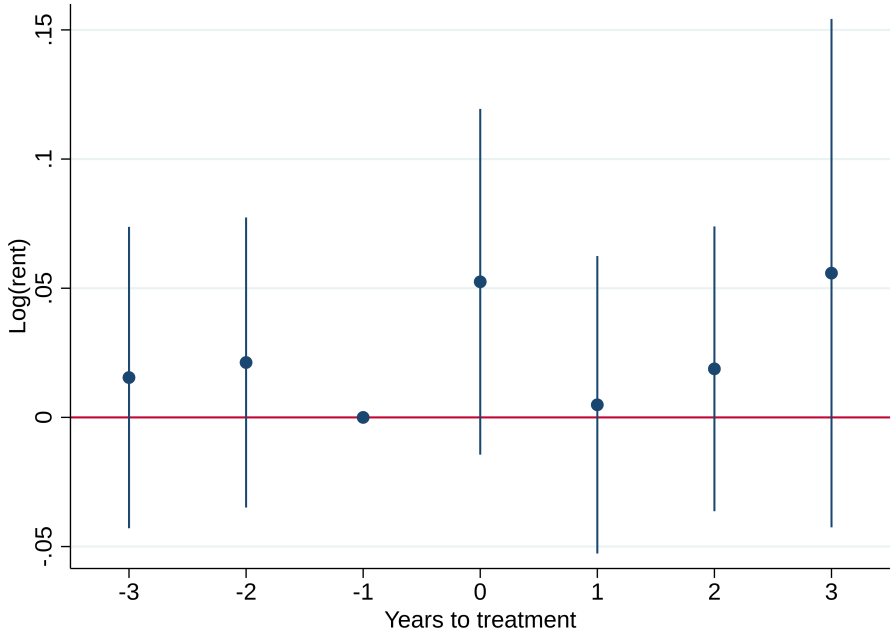

Note: This figure shows near-far results for the set of all new buildings. Panel A repeats the baseline near-far event study shown in Figure 3, while Panel B repeats the larger treatment radius robustness check shown in Figure A.6

Figure A.15: Near-Near Event Studies for Rent Outcome (All New Buildings)

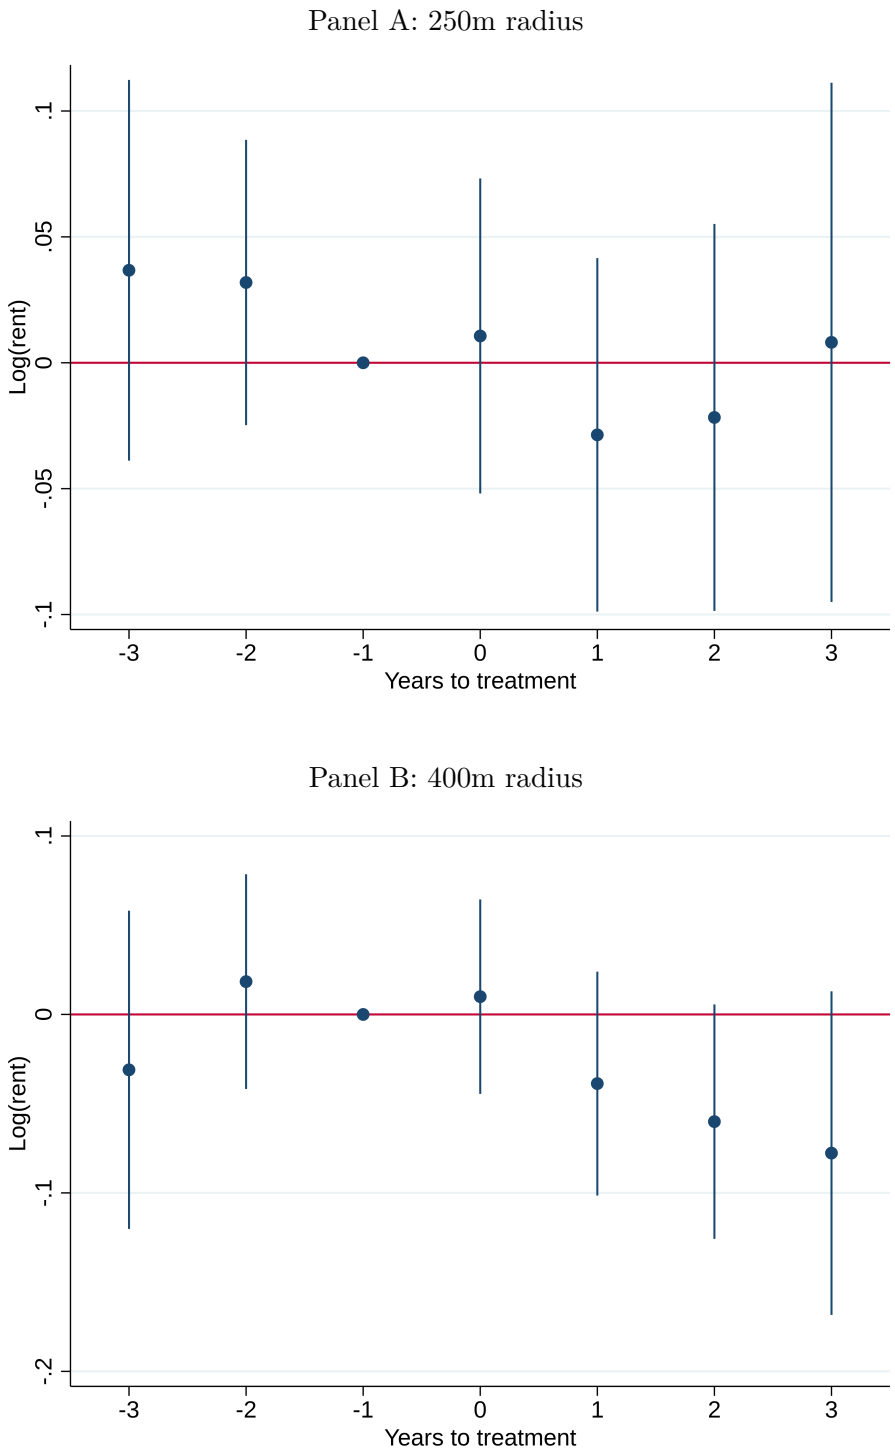

Note: This figure shows near-near results for the set of all new buildings. Panel A repeats the baseline near-far event study shown in Figure 4, while Panel B repeats the larger treatment radius robustness check shown in Figure A.7.

Figure A.16: Near-Near Event Study Robustness (Log(Origin Income))

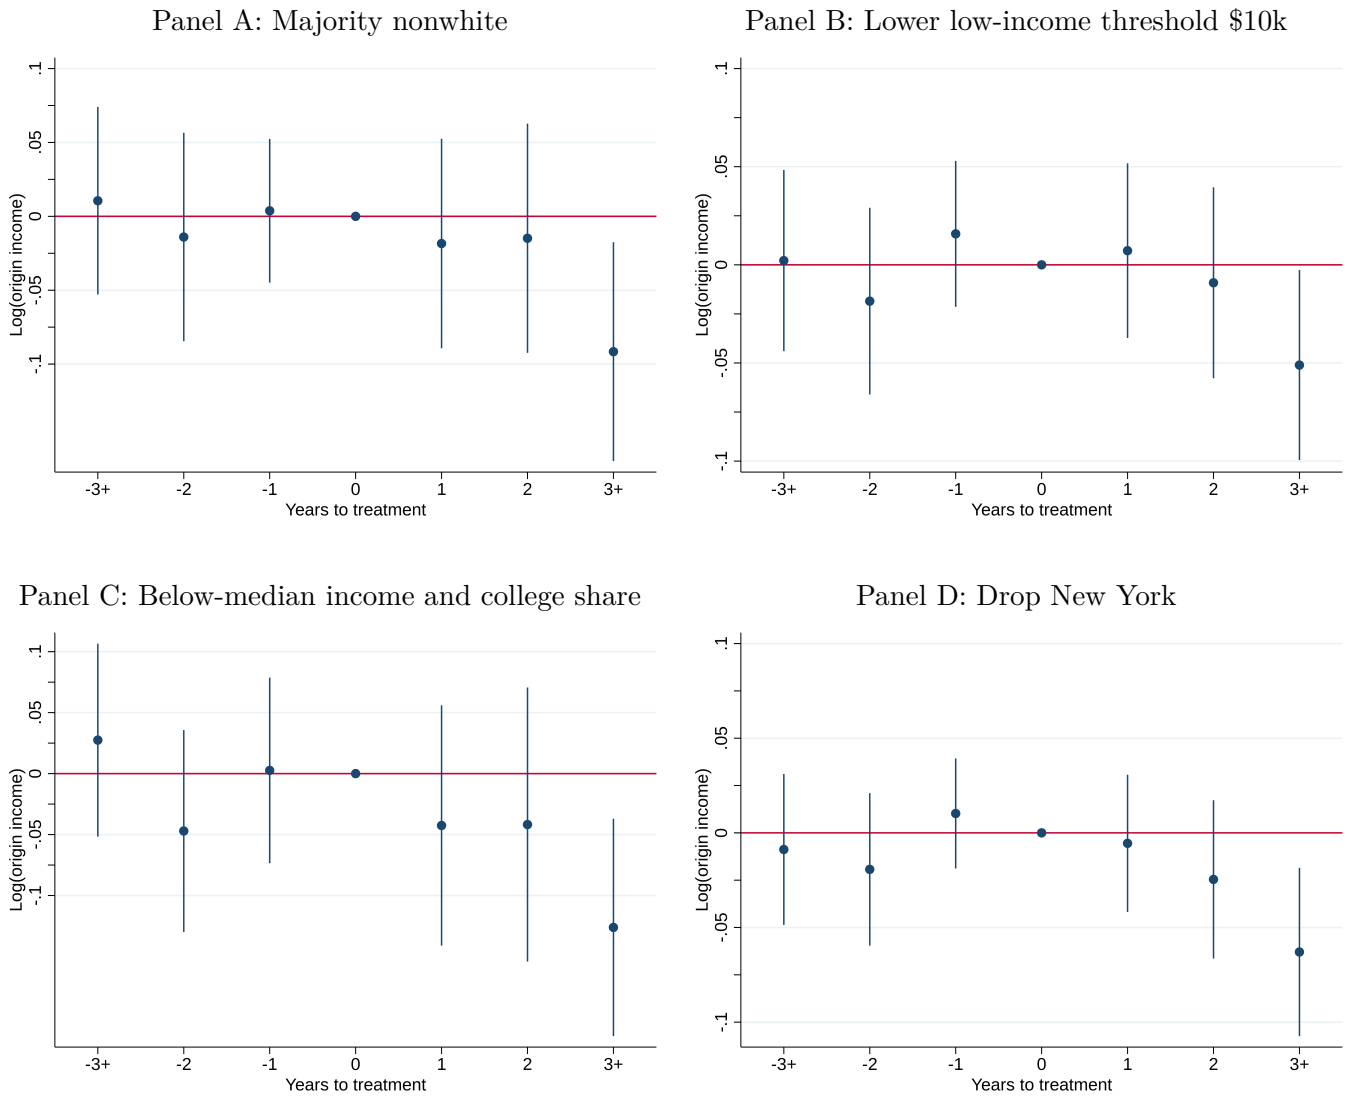

Note: Each panel repeats the baseline near-near event study for log(origin income) shown in Panel A of Figure 5 with a change to the sample. Panel A drops buildings in tracts that are over 50% white, and Panel B lowers the income threshold by $10,000. Panel C requires that both tract income and college share be below the CBSA median, while Panel D drops observations in New York City.

Figure A.17: Near-Near Event Study Robustness (High- and Low-Income Origins)

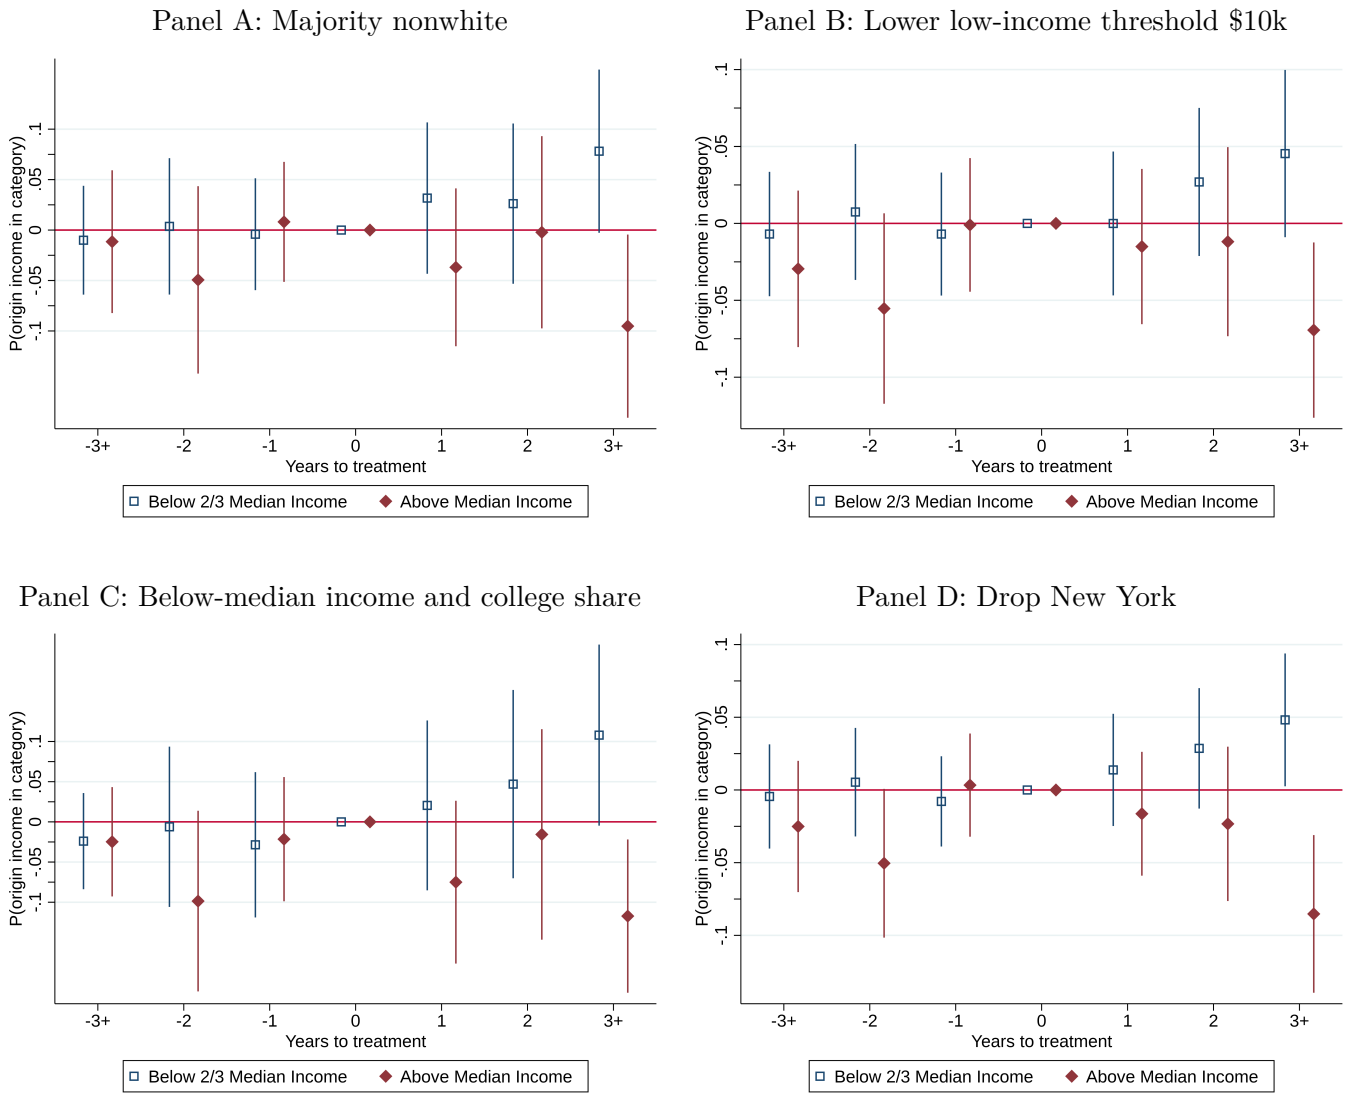

Note: Each panel repeats the baseline near-near event study for high- and low-income arrivals shown in Panel B of Figure 5 with a change to the sample. Panel A drops buildings in tracts that are over 50 percent white, and Panel B lowers the income threshold by $10,000. Panel C requires that both tract income and college share be below the CBSA median, while Panel D drops observations in New York City.

Figure A.18: Near-Near Event Studies for Migration Outcomes (400m Treatment Radius)

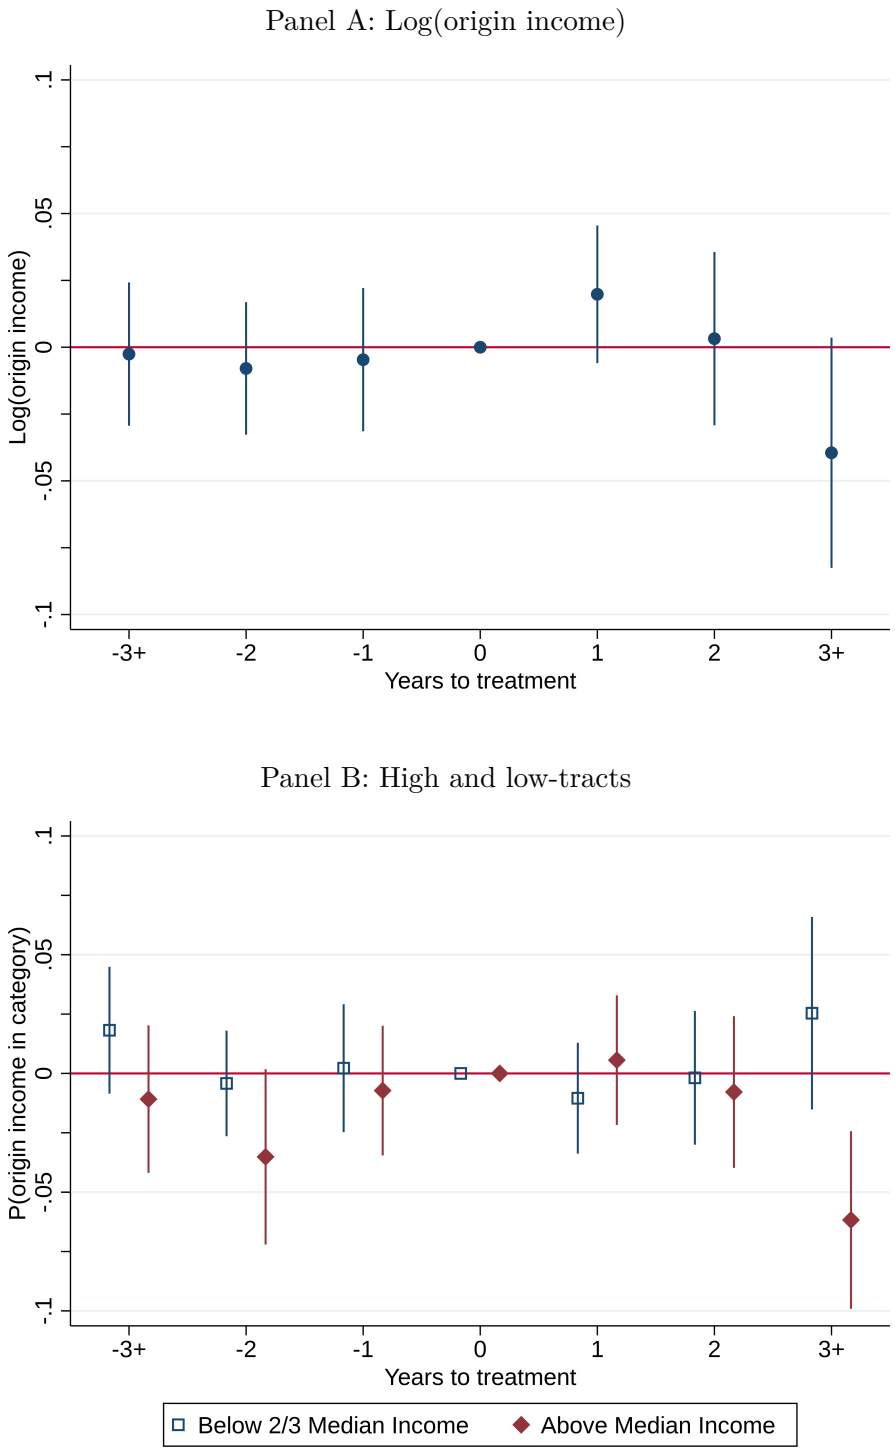

Note: This figure repeats the near-near event study shown in Figure 5 but increases the treatment radius from 250 meters to 400.

Table A.1: Mean Rents by Year and CBSA

	2013	*2014*	*2015*	*2016*	*2017*	*2018*
Atlanta, GA	1113	1174	1190	1314	1375	1477
Austin, TX	1309	1504	1581	1618	1636	1607
Chicago, IL	1680	1657	1947	1755	1680	1814
Denver, CO	1508	1688	1718	1730	1703	1762
Los Angeles, CA	1921	1965	2218	2431	2411	2309
New York, NY	3167	2692	2815	2849	2536	2187
Philadelphia, PA	1328	1235	1410	1399	1433	1458
Portland, OR	1339	1380	1513	1617	1595	1618
San Francisco, CA	2472	3103	3479	3621	3304	3304
Seattle, WA	1522	1581	1677	1847	1912	1912
Washington, DC	1958	2056	2073	2172	2181	2182

Note: This table shows the mean rent in a given year and CBSA in the provided sample of Zillow™ listings. The sample only includes listings in buildings with fewer than 50 units.

Table A.2: Mean One-Bedroom Rents by Year and CBSA

	2013	*2014*	*2015*	*2016*	*2017*	*2018*
Atlanta, GA	981	1045	1032	1234	1261	1402
Austin, TX	1005	1092	1139	1093	1149	1229
Chicago, IL	1481	1422	1749	1525	1469	1584
Denver, CO	1190	1385	1207	1226	1215	1299
Los Angeles, CA	1406	1344	1581	1761	1787	1780
New York, NY	2661	2343	2300	2363	1843	1741
Philadelphia, PA	1147	990	986	1039	1123	1196
Portland, OR	1079	1081	1224	1257	1234	1255
San Francisco, CA	2157	2736	2911	2922	2685	2693
Seattle, WA	1245	1307	1338	1467	1536	1541
Washington, DC	1682	1885	1736	1875	1794	1869

Note: This table shows the mean rent for one-bedroom units in a given year and CBSA in the provided sample of Zillow™ listings. The sample only includes listings in buildings with fewer than 50 units.

Table A.3: Rent Difference-in-Differences Results with 400m Treated Radius

	Near versus far	Near versus near	Triple-difference
After*within 400	-0.023		
(S.E.)	(0.025)		
After*treated building		-0.027	
(S.E.)		(0.032)	
After*within 400*treated building			-0.068
(S.E.)			(.041)
Treated buildings	47	44	43
Control buildings		20	15
Listing observations	41,500	28,400	46,400

Note: This table shows difference-in-differences results for rents with a 400 meter treatment radius. The first column shows the near-far specification shown in Equation 2, where the treatment group is listings within 400m of a building completed in 2015–2016, and the control group is listings between 400m and 600m of the same buildings. The second column shows the near-near specification, in which the treatment group is listings within 400m of a building completed in 2015–2016, and the control group is listings within 400m of buildings completed in 2019 (after the sample period). The third column shows the triple-difference specification from Equation 5, which compares the near-far gap in the 2015–2016 and 2019 buildings. Specification details are identical to Table 4, except the isolation restriction on new buildings is increased to 400 meters to match the treatment radius.

Table A.4: Far-Far Difference-in-Differences (Rent Outcome)

	250-600m band	250-800m band
After*treated building	0.028	0.015
(S.E.)	(0.026)	(0.021)
Treated buildings	73	80
Control buildings	28	31
Listing observations	41,000	60,900

Note: This table repeats the near-near rent specification, but compares the 250–600 meter (or 250–800 meter) bands around the treatment and control of buildings.

Table A.5: Near-Far Difference-in Differences Robustness (Rent Outcome)

	Low College	Low Percent White	Lower Income	No New York
After*within 250	-0.051	-0.074	-0.052	-0.061
(S.E.)	0.021	0.025	0.025	0.027
Treated buildings	21	25	37	36
Listing observations	30,308	35,840	52,799	24,178

Note: Each column repeats the baseline near-far DiD shown in Equation 2 with a change to the sample. Panel A drops buildings in tracts that are over 50 percent white, and Panel B lowers the income threshold by $10,000. Panel C requires that both tract income and college share be below the CBSA median, while Panel D drops observations in New York City.

Table A.6: Near-Near Difference-in Differences Robustness (Rent Outcome)

	Low College	Low Percent White	Lower Income	No New York
After*treated building	-0.16	-0.13	-0.069	-0.058
(S.E.)	0.042	0.05	0.037	0.043
Treated buildings	19	22	34	34
Control buildings	9	13	15	18
Listing observations	10,388	12,701	19,240	8,491

Note: Each column repeats the baseline near-near DiD shown in Equation 4 with a change to the sample. Panel A drops buildings in tracts that are over 50 percent white, and Panel B lowers the income threshold by $10,000. Panel C requires that both tract income and college share be below the CBSA median, while Panel D drops observations in New York City.

Table A.7: All-Income Difference-in-Differences Results (Rent Outcome)

	Near versus far	Near versus near	Triple-difference
After*within 250	0.016		
(S.E.)	(0.016)		
After*treated building		-0.017	
(S.E.)		(0.029)	
After*within 250*treated building			0.019
(S.E.)			(0.04)
Treated buildings	103	94	91
Control buildings		31	30
Listing observations	190,020	75,710	201,818

Note: This table shows difference-in-differences results for rents including all new buildings (instead of only those in low-income tracts). The estimations are otherwise identical to Table 4.

Table A.8: All-Income Difference-in-Differences Results, (Rent Outcome, 400m Radius)

	Near versus far	Near versus near	Triple-difference
After*within 250	-0.013		
(S.E.)	(0.018)		
After*treated building		-0.015	
(S.E.)		(0.023)	
After*within 250*treated building			-0.069
(S.E.)			(.033)
Treated buildings	95	95	85
Control buildings		23	20
Listing observations	120,032	80,320	127,292

Note: This table shows difference-in-differences results for rents including all new buildings (instead of only those in low-income tracts) and using a 400 meter treatment radius. The isolation restriction on new buildings is also increased to 400 meters to match the larger treatment radius. The estimations are otherwise identical to Table 4.

Table A.9: Far-Far Difference-in-Differences (Log(origin income))

	250-600m band	250-800m band
After*treated building	-0.01	-0.013
(S.E.)	(0.01)	(0.009)
Treated buildings	70	70
Control buildings	83	83
Listing observations	135,600	207,650

Note: This table repeats the near-near estimation for log(origin income), but compares the 250–600 meter (or 250–800 meter) bands around the treatment and control buildings.