

Local Effects of Large New Apartment Buildings in Low-Income Areas*

Brian J. Asquith Evan Mast Davin Reed[†]

April 2, 2021

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 rents in nearby units by about 6 percent relative to units slightly farther away or near sites developed later, and they increase in-migration from low-income areas. We show that new buildings absorb many high-income households and increase the local housing stock substantially. If buildings improve nearby amenities, the effect is not large enough to increase rents. Amenity improvements could be limited because most buildings go into already-changing neighborhoods, or buildings could create disamenities such as congestion.

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, Jenny Schuetz, ZillowTM, and Real Capital Analytics. Shane Reed, Nathan Sotherland, and Steve Yesiltepe provided excellent research assistance. 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. An earlier version of this paper circulated as "Supply Shock Versus Demand Shock: The Local Effects of New Housing in Low-Income Areas."

[†]Asquith: W.E. Upjohn Institute for Employment Research; Mast: W.E. Upjohn Institute for Employment Research; Reed: Federal Reserve Bank of Philadelphia.

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 and offset the effects of 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 represent a large change from the status quo (Hankinson 2018; Been, Ellen, and O’Regan 2019). Although this idea plays 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.

This paper provides new empirical evidence on how the construction of large market-rate rental apartment buildings in low-income, central city neighborhoods affects rents in nearby

¹Guerrieri, Hartley, and Hurst (2013); Diamond (2016); and Su (2019) study how endogenous amenities and residential sorting affect housing costs. Baum-Snow and Marion (2009) and Diamond and McQuade (2019) show that building low-income housing can increase nearby home prices in some cases.

buildings 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.² We find that new buildings lower rents in nearby buildings by 5 to 7 percent relative to trend and increase in-migration from low-income areas. If a positive endogenous amenity effect exists, it appears to be overwhelmed by a standard supply effect. Alternatively, buildings could instead create disamenities like congestion. We provide evidence that supply is a major driver of rent decreases but cannot rule out that disamenities contribute. 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 the regional and local price effects of new housing construction.

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). 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 positive amenity effects that fade out quickly with distance, which prior literature has shown to be a common shape of housing externalities.³

The second exercise is a difference-in-differences that compares listings near buildings

²Our sample includes Atlanta, Austin, Chicago, Denver, Los Angeles, New York City, Philadelphia, Portland, San Francisco, Seattle, and Washington, D.C.

³Our primary specification uses 250 meters as the treatment group and 250–600 meters as the control.

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. This directly controls for characteristics that make developers target certain parcels. Moreover, 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. 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 ZillowTM that span 2013 to 2018. We focus on “pioneer” buildings—those that are the first recent market-rate construction in their area—because these are most likely to change a neighborhood and because buildings subsequently constructed nearby may be part of the pioneer building’s effect. All three empirical approaches show that new buildings in low-income neighborhoods (census tracts) reduce nearby rents by 5 to 7 percent.⁴ 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 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. In extensions, we show that changes in the composition of units listed on Zillow do not explain our results and that other possible spatial patterns of rent increases, such as a “doughnut” caused by congestion effects near a

⁴We define low-income neighborhoods as census tracts with median household incomes below the metropolitan area median. We include all buildings 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 higher than for gentrifying low-income areas, reducing a supply shock’s effect on prices.

new building, are unlikely.

In our second set of results, we study the effect on in-migration using individual address histories from Infutor Data Solutions. In-migration allows us to study cheaper segments of the market that may be underrepresented in the ZillowTM data. In addition, while popular discussion of gentrification often focuses on out-migration, prior research shows that neighborhood change actually occurs primarily through shifts in the quantity and composition of in-migrants, making this the more policy-relevant outcome (Freeman and Braconi 2004; McKinnish, Walsh, and White 2010; Brummet and Reed 2019). Nonetheless, we present some summary statistics on out-migration in Section 1.3, but data limitations make it difficult to study this outcome within our DiD framework.

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 two 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. Triple-difference results are similar, while the near-far specification returns a null result.

Finally, we consider the supply and amenity effects underlying our main results. Given that the average sample building increases housing supply within 250 meters by 15%, they should put downward pressure on prices. However, it is less obvious that supply can produce such a local effect. A few results from prior literature provide support. Anenberg and Kung (2014) find that the competitive effect (net of disamenities) of a foreclosure listing creates a similarly steep price gradient, and both survey evidence and online housing search data show that renters are highly sensitive to location. We also use a simple geometric exercise to illustrate why added supply can have very different competitive effects across short distances even when searchers consider larger areas. On the amenity side, positive effects may be small because, as shown in Section 1.1, new construction typically occurs after a neighborhood has already begun to change. Even among low-income neighborhoods, those that received

new buildings saw more income and education growth over the previous decade, which may limit the additional signal or amenity changes that a building can create. Alternatively, the new building could create disamenities that reinforce the negative supply effect. We cannot fully disentangle these two forces, but the negative price effect is crucial to the policy debate regardless of the underlying mechanism.

These results matter for policy. Approving new housing in low-income areas is often contentious because of worries that new buildings will accelerate nearby rent increases. These local concerns also spill over into the regional policy debate and can stall large-scale housing reforms (Brey 2019; Dillon 2019). While new buildings may increase average income in the area (or even average rents *inclusive* of the new units), our results suggest that they do not increase rents in the existing housing stock. Policymakers should recognize that new market-rate housing can improve affordability both locally and regionally, making it an important component of strategies to address the growing affordability crisis. In addition, our migration results suggest that strategies that encourage housing construction may 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 apartment buildings 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. We also do not study housing or land prices, which depend on future expectations in more complicated ways than rents. 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 growing literature on the effects of new housing. One branch focuses

on regional effects and takes a model-based approach, generally finding price reductions (Anenberg and Kung 2018, Mast 2019, Nathanson 2019). In contrast, we study very local effects, which may be substantially different. Li (2019) also studies local effects, finding that large new buildings in New York City lower nearby buildings' rental income even as they increase the number of nearby restaurants. Finally, Damiano and Frenier (2020) use data on rent in large buildings in Minneapolis and find no average effect of new construction, but a 6.7% increase in bottom-tercile rent buildings and 1.7% decrease in the highest tercile. We differ from these papers by studying many cities, using listing-level rents, and examining migration outcomes.

More broadly, a large literature considers the effects of land-use regulation (Gyourko and Molloy 2015, Glaeser and Gyourko 2018). It shows that restrictive regulations lead to higher rent and price growth (Pollakowski and Wachter 1990; Quigley and Raphael 2004; 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; Buntin 2017; Parkhomenko 2018). Our work adds to this literature by suggesting that building housing improves neighborhood as well as regional affordability.

The body of the paper is organized as follows. Section 1 describes the data and provides summary statistics, and Section 2 describes our empirical strategy. Sections 3 and 4 present our rent and migration results, respectively, and Section 5 concludes.

1 Data and Summary Statistics

1.1 New Buildings Data

Data on large new market-rate rental apartment buildings are provided directly by Real Capital Analytics, a real estate market research firm. 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.⁵ These are large, supply-constrained cities that have experienced substantial reurbanization since 2000 and anecdotally have heated debates over gentrification and new market-rate construction. The data include year completed, number of units, and exact address, which we match to 2010 census tracts (our definition of a neighborhood). In addition, there is a flag for features like income or age restrictions, helping to identify, for example, buildings financed under the LIHTC program.

However, not all of these buildings are relevant to our research question. The policy debate focuses on low-income areas and buildings that represent large changes to their neighborhood, where the mechanisms that could spark nearby rent increases are likely stronger. We apply the following restrictions to construct our analysis sample.

1. *Large market-rate rental buildings:* We remove buildings that RCA flags as income-restricted, senior, or student housing, as well as those in tracts with over 25% college students. We also restrict to buildings with 50 units or more, which are most contentious and most likely to change a neighborhood physically or demographically.
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 (ACS).⁶ 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. *“Pioneer” buildings:* We include only buildings that had no other new buildings

⁵We restrict to the borough of Brooklyn when studying New York because this is the only borough included in our data use agreement with Zillow™.

⁶We 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.

completed within 250 meters (the baseline treatment radius) between 2010 and their date of completion. Note that this does *not* exclude buildings where more buildings were later constructed nearby. This restriction identifies buildings that represent a larger shock to a neighborhood. In addition, it prevents us from overcontrolling for the effects of the first building in a neighborhood—subsequent construction nearby may be part of the initial building’s effect. However, this restriction does preclude estimating the effect of additional buildings later constructed near a pioneer building.

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.⁷ This improves the precision of our estimates and helps ensure we are picking up buildings in preexisting neighborhoods where gentrification may be a concern, rather than large brownfield redevelopments.
5. *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 we have a sufficient number of years before and after treatment in our outcome data.

[TABLE 1 HERE]

Table 1 shows building characteristics and location for four subsamples. The first column, “All Incomes 2010–2019,” restricts to large, market-rate buildings in central cities completed between 2010–2019. The second column further restricts to buildings built in low-income and non-student neighborhoods, yielding 768 buildings. The third drops buildings outside of the years of our analysis sample, leaving 231. Finally, the last column adds the sufficient listings restriction and includes only pioneer buildings, yielding the final set of the 96 buildings that

⁷The 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 250m and 600m. For the near-near approach, the treatment area is within 250m of early buildings, and the control is the same distance from later buildings.

are included in the main specification for either rent or migration outcomes. The small final set highlights that pioneer buildings in low-income residential areas are a small proportion of all construction, underscoring the importance of a multi-city sample for statistical power.

Despite the loss of observations, Table 1 suggests that this group of buildings is representative of the full set of low-income buildings in the RCA data. The median number of units in the analysis sample is 117, versus 141 for all buildings in low-income areas. For the median building, this represents a large 37 percent shock relative to the existing stock of 320 rental units within 250 meters (the baseline treatment radius) and a much smaller 6 percent shock to the 1,843 rental units within 600 meters (the baseline control radius).⁸ In addition, the median ZillowTM listed rent within 600 meters for the median building in the final sample is \$1,836, similar to the same figure for all low-income buildings completed in 2014-2016.

We next consider the types of neighborhoods new developments are built in. Table 2 describes characteristics of central city neighborhoods 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, and columns 3 to 6 represent different groups of low-income neighborhoods. The rows show levels and growth rates for income, college education, and rents.

[TABLE 2 HERE]

In both the all-income and low-income samples, neighborhoods with construction experienced much faster preperiod growth in income and college share and somewhat faster growth in rents. While the levels of income were similar across low-income samples during the preperiod, the areas that received construction had higher college shares, which is often considered a leading indicator of gentrification. Overall, the patterns in Table 2 suggest that

⁸We estimate the number of units within these radii based on the average density of rental housing units in each building's tract in the 2013-2017 ACS. Within both radii, about 80 percent of the stock is rentals.

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, unlike low-income areas that are not gentrifying, they are already attractive to high-income residents who can pay the higher rents required for new construction.

Finally, the last column shows the neighborhood characteristics for our final sample of analysis buildings, or “pioneer” buildings. Three points are worth emphasizing. First, these buildings, although in areas without any recent construction, still appear to be in neighborhoods that are already gentrifying. This is our strongest evidence that new buildings typically follow incipient gentrification, rather than initiating the process. Second, the neighborhoods containing pioneer buildings do not appear to be different from low-income neighborhoods that receive new buildings in general. Finally, 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.

1.2 Rental Data

We pair the new building data with listing-level data on rental prices provided by ZillowTM, which includes listings from all websites (ZillowTM, Trulia, StreetEasy, and HotPads) in the ZillowTM Group during 2013 to 2018. For each listing, we observe price, location, date posted, and the number of bedrooms and bathrooms.⁹ ZillowTM obtains the listings in our sample directly from landlords who advertise rentals on the platform. The major restriction is that it includes only buildings with 50 units or less, which is the subsample that ZillowTM, for internal reasons, was able to share. The advantage of this sample is that there are many small buildings, providing the dense coverage needed for spatial identification strategies.

In total, we observe rents for about 740,000 units within 800 meters of one of our new buildings. Appendix Figure A.1 shows the listings within 800 meters of a new building in

⁹Due to a few extreme values, we winsorize rents at the 1st and 99th percentile.

Chicago in 2018, and 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.

Although online search is very common—a 2016 ZillowTM survey found that 84 percent of renters use online resources—the listings in our sample may not be perfectly representative (Curnette et al. 2016). We assess this by comparing tract-level median rents in our sample of ZillowTM data and the 2013-2017 ACS. To obtain the closest possible comparison, we calculate the ACS median contract rent for households who moved into their unit between 2015 and 2017 and ZillowTM listings for the same years.¹⁰ Across all low-income tracts in our sample, the average median rent in ZillowTM is \$1,590 for a one-bedroom and \$2,040 for a two-bedroom, which is around 20 percent higher than the corresponding ACS rents of \$1,320 and \$1,720.

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 sources—USPS change of addresses, county assessor records, magazine subscriptions, phone-books, et cetera—and sells it 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, along with 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

¹⁰We calculate the ratio of rents paid by households who moved in during 2015-2017 to rents paid by all households in each tract. We then multiply the ratio by the median one- and two-bedroom rents in the tract (which are not available by year moved in). Finally, we subtract median tract-level utility costs.

across tract characteristics.¹¹ 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.¹²

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. We focus on this in-migration measure, instead of more frequently discussed out-migration metrics, primarily because the academic literature has reached a strong consensus that neighborhood change primarily occurs through changes in in-migration.¹³ In-migration also allows us to infer something about rent in listed units. In addition, data limitations make it difficult to study out-migration. First, directly measuring displacement requires identifying incumbent individuals in the area and then following their movement longitudinally. Since Infutor misses some moves, a longitudinal address history of an individual will likely contain more measurement error than our in-migration outcomes, which only require the start and end points of one move. Moreover, limited individual characteristics in the Infutor data mean we cannot separate low- and high-income individuals within the same area, making it impossible to identify low-income out-migrants.

[FIGURE 1 HERE]

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

¹¹Appendix Figure A.2 plots the ratio of the Infutor and Census populations against tract characteristics. Diamond et al. (2019) and Phillips (2019) provide additional validation.

¹²We assess this by comparing county-level migration rates in the Infutor data to census estimates. Appendix Figure A.3 plots the ratio of the two estimates against county characteristics and shows only slight correlations, suggesting that observed moves are close to randomly selected.

¹³See Freeman and Braconi (2004), McKinnish, Walsh, and White (2010), Ding, Hwang, and Divringi (2014), and Brummet and Reed (2019).

trends in in-migration within 250 meters of the 2014–2015 buildings in our low-income analysis sample. The triangle-marked line shows the average number of migrants to a new building, the square-marked line shows the average number of migrants to the surrounding 250 meters excluding moves into the 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. 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 show in Appendix Figure A.4 that net migration from low-income areas does not 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. In order to maintain a large sample, we use spatial strategies that can be applied in all cities, rather than using policy-driven variation in a single city.

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 effects are large, we could see rent increases very close to the new building that then 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 (2018) and Autor et al. (2014) take a comparable approach in other settings. More generally, Brooks and Lutz (2016) highlight the difficulty of assembling large parcels of land.

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 groups. The idea is that within a small area, developers have few sites that are available and properly zoned, 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 unlike the treatment area. However, there is a tradeoff, as the treatment area must be large enough to be substantially more affected by the new building than the control area. Prior literature suggests that the effects of both new housing supply and related externalities decay very quickly.¹⁴ This motivates our choice of a baseline treatment radius of 250 meters (roughly one or two city blocks) and control radius of 600 meters (slightly less than half a mile, an 8- to 10-minute walk).¹⁵ We investigate the validity of the identification assumption by estimating an event study specification 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

¹⁴Anenberg and Kung (2014) use a 160-meter radius to detect the supply effect of new foreclosure listings, while Diamond and McQuade (2019) show similarly quick decay in the price effect of new LIHTC buildings. Schwartz et al. (2006) and Rossi-Hansberg et al. (2010) find that home price effects of urban revitalization programs decay substantially within 250m.

¹⁵Time estimates are from Google Maps and vary depending on the road network.

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 κ_{br} at the level of the nearest building \times treatment status. We also use the nearest building to define time fixed effects α_{bt} that control for time-varying shocks at a very local level.

This yields the following specification for rent in listing i in year t with treatment status r and nearest building b :

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

X_{it} contains dummies for the number of bedrooms and bathrooms, controlling for potential 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:

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

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 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 the idiosyncratic nature of the land assembly, building permitting, and construction processes.

We evaluate this assumption in Appendix Table A.3 by regressing pre-period (2010-2014 ACS) characteristics of each building's census tract on CBSA fixed effects and a treatment dummy. Because we later repeat the near-near strategy with migration outcomes, we include buildings in both the rent and migration analysis samples. Estimates are generally small and statistically insignificant at conventional levels—treatment and control groups appear to be at very similar stages of gentrification in the years prior to the treatment buildings' completion. Only one of 18 coefficients is statistically significant at a 10 percent level (the 2010-2014 change in the share of residents with a college degree). None of the other common proxies of gentrification, including levels and changes in incomes, rents, house values, and population, are statistically or economically significant. This suggests that developers face frictions that prevent them from precisely timing new buildings. Within our small time window, treatment timing instead appears to be driven by idiosyncratic variation, consistent with our identification assumption. We later probe this assumption further by examining pretrends in event studies.

To implement the near-near 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 include time fixed effects α_{mt} at the CBSA \times year level (rather than the nearest building level as in the near-far). Finally, we include location fixed effects κ_b at the nearest building level. 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.¹⁶

This yields the following event study specification for listing i in year t in CBSA m with treatment status c and nearest building b :

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

We weight observations, cluster standard errors and define X_{it} as in the near-far specification.

We also estimate the standard DiD:

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

While both the building's announcement and completion could be important for the nearby housing market, completion is likely more relevant for rents. Amenity effects depend on the building directly changing the area or its residents attracting new businesses or people, neither of which occurs until completion. While expected future amenities may change with a building's announcement, given the short average tenure of renters and typical length of leases, it is unlikely rents will be affected by improvements that will not occur for years.

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

¹⁶Because 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.

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 treatment and control inner rings (such as the treatment group gentrifying earlier), so long as the difference is the same in the two outer rings.

To implement the triple-difference specification, we must first expand the sample to include all listings within 600 meters of any 2015, 2016, or 2019 building. Note 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). This yields the following specification for rent in listing i in year t with near-far treatment status r , near-near treatment status c , and nearest building b :

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

where ν_{rt} is a fixed effect for the inner ring in each year and α_{bt} and κ_{br} are the fixed effects from previous specifications. We use the same weighting and clustering schemes as in the previous specifications and again include bedroom and bathroom counts in X_{it} .

3 Rent Results

3.1 Main Results

Results for the near-far specification are in Figure 2 and Table 3, Column 1. The event study specification in Figure 2 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 3, Column 1) shows that new buildings decrease nearby rents by 4.9 percent (S.E.=0.021, $p=0.023$). These results suggest that highly local positive spillovers do not cause rents to rise in the immediate area of the building.

[FIGURE 2 HERE]

Results for the near-near specification, which allows for broader effects of new buildings, are shown in Figure 2 and Table 3, Column 2. They again 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 -0.070 following building completion. The after-treatment estimate from a standard DiD is -0.062 (S.E.=0.037, $p=0.096$). This is about 25 percent larger than the near-far estimate (though the difference is not statistically significant), perhaps because spillover effects on the control group attenuate the near-far estimate. Finally, the triple-difference DiD estimate appears in Table 3, Column 3 and is quite similar: -0.071 (S.E.=0.033, $p=0.037$).

[TABLE 3 HERE]

The results from above, as well as from a number of robustness checks, are summarized together in Figure 3. In dollar terms, the point estimates from the main specifications translate to between a \$100 and \$159 decrease in listed 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.

[FIGURE 3 HERE]

The baseline estimates use a treatment radius of 250 meters, which we extend to 400 meters in Appendix Table A.5.¹⁷ The near-near estimate decreases to a statistically insignificant

¹⁷The radius in the pioneer restriction is extended to 400m to match the treatment radius.

−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 Figure A.6) 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 suggests 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 mechanisms and robustness, we discuss two caveats that we cannot explore 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—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 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, as discussed in Section 1.2, the data from ZillowTM 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 initially lower-rent units. While we have no other source of listing-level rent data to study this directly, our migration data allow us to consider the lower end of the housing market by examining migration from low-income areas.

3.2 Discussion and Mechanisms

Our near-far and near-near strategies are designed to test for perhaps the most likely way that new buildings could increase rent: through endogenous amenity improvements centered on the new building. However, the negative estimates instead point to the importance of

supply effects and possible disamenities. This may be because the low-income neighborhoods that see new market-rate construction tend to already be gentrifying, limiting amenity improvements. In this section, we discuss the contribution of supply and disamenities to our main results.

Increased housing supply should, of course, lower nearby rent. Given that these buildings are large relative to the existing stock—Table 1 shows they represent a 40 percent increase in rental units within the treatment area—it is plausible that they move prices a measurable amount. However, our results show that the negative effect is highly concentrated, fading out substantially within a couple blocks. It is intuitive that the effect should be largest close to the building but less obvious that increased supply can generate such a steep gradient.

We provide two pieces of evidence that a supply effect can create this pattern. First, there are a variety of supporting findings in the prior literature. Most closely related is Anenberg and Kung (2014), who find that the competitive effect (net of disamenities) of a listed foreclosed home is nearly twice as large within 160 meters as between 160 and 533 meters. Other findings suggest that renters are very location-sensitive, which is consistent with hyper-local supply effects. Rae and Sener (2016) examine user-drawn search areas on a major online housing platform in London. Twenty-three percent span less than a square kilometer, corresponding to a circular choice set radius of about 550 meters. Moreover, people searching for smaller units and higher prices—likely the relevant group—search in smaller areas. Outside of the academic literature, an Apartment Guide survey found that only 14 percent of renters are willing to compromise on location, and a 2017 Zillow report found that 75 percent of renters ultimately leased a unit in the area where they began their search (Sirull 2018; Zillow Group 2017). These surveys are vague on the meaning of “area” or “location,” but they likely translate to a small area in our sample of dense cities.

Second, we use a simple geometric exercise to illustrate why the price gradient could be steep. Suppose that renters consider apartments within a circular choice set with radius r and a center that is randomly drawn with uniform probability across space. Given a distance

d between a unit and the new building, it is easy to compute the probability that a renter considering that unit also has the new building in their choice set.¹⁸ While this simple metric of competition does not describe equilibrium price responses, it does illustrate why important differences might exist over short distances. With $r = 550$ and $d = 200$, a unit will compete with the new building for 77% of its prospective renters. At $d = 500$, the probability falls to 44%. With $r = 800$, implying that a renter will consider units a full mile apart, these probabilities are 84% and 61%. Even though the example units are only 300 meters apart, there is a substantial difference in the intensity of competition from the new building.¹⁹ Appendix Figure A.7 shows the metric for d between 0 and 1000 for r of 550 and 800 meters.

This evidence suggests that a very local supply effect is likely an important driver of our estimates, but it does not rule out disamenity effects. New buildings may create congestion, block views or light, or be aesthetically unappealing. Given our focus on rents, it is unnecessary to fully distinguish between negative amenity and supply effects, though local disamenities would of course be a part of a comprehensive welfare assessment (along with the effect on rents elsewhere and many other factors). However, a problem could arise if negative amenity effects near the building drive our main estimates but positive amenity effects slightly farther away actually do increase rents. 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 positive rent effect in a doughnut shape.

¹⁸Since the unit is in the renter's choice set, the set's center is within r meters of the unit. If the center is within r of the new building, the renter will also consider it. The probability of interest is then simply the area of overlap between two circles with radius r and centers d meters apart divided by πr^2 . The area of overlap is given by a standard geometric formula.

¹⁹The fact that both treatment and control group face some competition illustrates the importance of our near-near approach and may explain why that estimate is 25 percent larger than the near-far.

The simplest way to test this story is to repeat the near-near strategy but compare the 250- to 600-meter bands in the treatment and control groups, excluding the inner area most affected by congestion. Appendix Table A.6 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.

However, this “far-far” approach checks for a very particular effect shape with a set size, and we do not know *ex ante* where rent increases might occur. To observe patterns more generally, we use Diamond and McQuade’s (2019) empirical derivative approach, which produces continuous estimates of the treatment effect on rents at various distances to the new buildings. The primary drawback is that it only accommodates buildings that had no other construction within the studied area during the entire sample period, which reduces our sample size and results in a potentially selected set of buildings.²⁰ Appendix Figure A.8 shows the estimated treatment effect at distances within 600 meters of the new building. 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 within our 250-meter treatment radius.²¹ Beyond 250 meters, the gradient rises toward zero as we approach the 600-meter limit. There are no nonlinearities that would suggest that the building increases rents in a doughnut or some other spatial pattern that would evade our primary specifications.

3.3 Robustness

Changes to Composition of Listings: Our results could be biased downward if landlords of lower-end units become more likely to list on Zillow in response to a new building, perhaps in

²⁰See Appendix A.1 for more detail on implementation.

²¹The estimated effect very near the building is –40 percent. This may be exaggerated because the number of listings shrinks with the square of the radius, leading to a very small sample and noisy estimates.

anticipation of more high-income people considering the area. We test for such compositional changes by looking at the effect of building completion on unit characteristics such as bedrooms and bathrooms and on the number of listings in treatment versus control areas. Appendix Table A.4 shows that bedroom and bathroom counts generally decline slightly after the new building goes in, but only one of the six coefficients is statistically significant at the 10% level. The exception is bedroom counts in the near-far specification, where the estimate is -0.10 (S.E. = 0.061 , $p=0.094$), though the coefficient changes signs across specifications and is of small economic importance. Total listings in the treatment area appear to decline slightly, although none of the results are statistically significant.

More Restrictive Neighborhood Definitions: Appendix Table A.7 shows near-far DiD results under three different restrictions—dropping majority white tracts, reducing the income threshold by \$10,000, and dropping listings in New York.²² In all cases, the patterns appear similar to the baseline, with estimates between -5 percent and -7 percent. Appendix Table A.8 shows results for the same changes to near-near specification, which again appear similar to the baseline.

Drop Pioneer Building Restriction: Appendix Figure A.9 shows near-far and near-near event studies after dropping the pioneer building restriction, which roughly doubles the number of buildings in each strategy. Strong negative pre-trends appear in both cases, likely because of the confounding effect of buildings that were completed nearby in the pre-period. While these event studies suggest that removing this restriction results in a less clean natural experiment, the corresponding difference-in-differences coefficients are not drastically different from our primary results. The near-far result is -0.030 ($t=1.88$) (Appendix Table A.7, Column 1), and the near-near is -0.033 ($t=1.4$) (Appendix Table A.8, Column 1).

²²The number of listings and average rent in New York change sharply in 2017, as shown in Appendix Table A.1. Although we include CBSA \times year or nearest building \times year fixed effects in all specifications, dropping these observation provides another check that this change does not drive our results.

Longer Time Periods: 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 post-construction years (Appendix Figure A.10), and find no evidence of long-run changes in the treatment effect. 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.11). We also study longer-run pretrends using the empirical derivative approach (Appendix Figure A.12) 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.

3.4 Including All Neighborhood Incomes

To conclude the rent section, we remove the restriction on neighborhood income and repeat our main specifications to examine whether effects may be different outside of low-income neighborhoods. Table 4 shows that DiD estimates with the 250-meter radius are close to zero and not statistically significant. Appendix Table A.9 shows that the estimates for the 400-meter radius are similar except for the triple difference, which is statistically significant and negative. The event studies, shown in Appendix Figure A.13 for the near-far specification and Appendix Figure A.14 for the near-near, are also noisy and inconclusive. These differences from the low-income results could occur because effects really are different in high-income areas or because our empirical strategy is less effective in such neighborhoods.

[TABLE 4 HERE]

There are several reasons that the true effect of new buildings could be different in

high-income areas. Amenity effects depend on the context, and the effect of adding supply depends on the underlying price elasticity of demand. One would expect positive amenity effects to be smaller in high-income areas that already have a large wealthy population and the associated reputation. However, because higher-income areas appeal to a broader swath of the population than gentrifying or low-income areas (Guerrieri, Hartley, and Hurst 2013), they may see a larger demand response to a given change in prices, muting the supply effect of new buildings.²³ This could occur because, for example, rich neighborhoods are perceived to have better schools, making them a plausible option for high-income families with children who would not consider a gentrifying area regardless of price. Real or perceived crime rates could have a similar effect.

In addition to substantive reasons for different effects, high-income real estate markets may be different in ways that matter for our empirical strategy. To examine this, we repeat the main specifications including only high-income buildings. While the near-near event study shows no consistent pattern, the near-far event study shows a strong upwards trend in both the pre- and post-period (Appendix Figure A.15).²⁴ One reason for this trend may be that large developable sites in high-income areas are less common and, given that they were not previously developed despite the area's affluence, are potentially strange in a way that affects nearby rents and contaminates the comparison between the inner and outer rings. Examples of such sites include formerly institutional uses like hospitals or public housing,

²³Anenberg and Kung (2018) and Glaeser and Ward (2009) both find that housing costs in expensive areas are relatively unresponsive to added supply and discuss how a high elasticity of demand could contribute to this result. In contrast, Busso et al. (2013) find a low population response to a large place-based policy targeting very low-income areas, suggesting that few people are on the margin of moving in.

²⁴While the event studies cast some doubt on causally interpreting results in these areas, the DiD coefficients corresponding to these event studies are 0.068 ($t=2.77$) for the near-far and 0.022 ($t=0.48$) for the near-near.

legacy manufacturers in urban cores, or relocated infrastructure like train tracks or a highway on-ramp. This explanation also explains why the trend does not appear in the near-near strategy, which compares similarly unusual sites that were developed at different times.

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 in the near-near specification. 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.²⁵

4.2 Migration Results

We begin with results from the near-near specification. Event study coefficients for log origin income are shown in Figure 4, Panel A. Pre-period coefficients are statistically indistinguishable from zero, suggesting parallel trends. This is followed by a sharp decrease of approximately 3–5 percent in the three years following building completion. The drop

²⁵For 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.

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, $p=0.085$).

[FIGURE 4 HERE]

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 cheaper segment of the housing market. To measure this effect, we repeat the specification using an indicator for whether an in-migrant is from a low-income tract (defined as below two-thirds of the CBSA median) as the dependent variable. Event study results are in Figure 4, Panel A, and they show a parallel pretrend that breaks sharply upward one year after treatment. The average effect, shown in Panel B of Table 5, is 0.034 (S.E. = 0.017, $p=0.059$). 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. These results also 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 4, Panel B, show no consistent pattern: across both 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 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 buildings tend to be constructed in areas already zoned for relatively high-density housing.

[TABLE 5 HERE]

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 Tables A.10 and A.11 show DiD results for log origin income and low-income arrivals under a number of tweaks to the sample: reducing the income threshold, restricting to less white neighborhoods, removing the pioneer building restriction, and dropping New York City. All are similar to the baseline results. In addition, event studies with a 400-meter radius are shown in Appendix Figure A.16. Similar to the rent results, these do not show the same sharp breaks as the baseline specification.

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.12, 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 (Been, Ellen, and O'Regan 2019; Hankinson 2018). However, there is little related empirical evidence. We fill this gap using quasi-experimental methods and address-level microdata.

We find that the concerns about rent increases driven by new market-rate housing are mostly unfounded. While there is a strong observed correlation between new construction and rising rents, this appears to be because new buildings are typically constructed in areas that are already changing. When these new buildings are completed, they actually slow rent increases 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 to \$159 per month. 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 improve economic mobility (Chetty, Hendren, and Katz 2016; Chetty et al. 2018).

We suggest that the mechanism underlying these results is a simple story of supply and demand. If high-income households like a particular neighborhood, preventing the construction of new housing simply leads them to outbid lower-income people for the housing already available in that neighborhood. By contrast, when new housing is built, many high-income households choose this option instead of nearby existing units. While the new

building could theoretically change local amenities or reputation by enough to instead increase demand and raise nearby rents, our findings suggest this is not the case. This may be because new buildings typically go into areas that are already changing, limiting their marginal reputational effect. Disamenity or congestion effects could also contribute to rent decreases near the building, but prior literature on hyper-local supply effects and spatially concentrated housing search suggest that added supply plays an important role.

Our results suggest that increasing housing supply can play an important role in addressing the present affordability crisis. Regulations that make it difficult to build could be relaxed, and localities could be directly incentivized to increase housing production. However, there are several caveats to our findings. First, our sample consists of areas where developers actually chose to, and were allowed to, build in. While these are likely similar to the neighborhoods that would receive new construction if regulations were relaxed, effects may be different in other types of neighborhoods. Second, relaxing land-use regulation is quite complicated in practice. 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. Finally, we consider only rents, not prices for land or existing homes. 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. (2014). Estimates of the size and source of price declines due to nearby foreclosures. *American Economic Review*, 104(8):2527–2551.
- 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*.

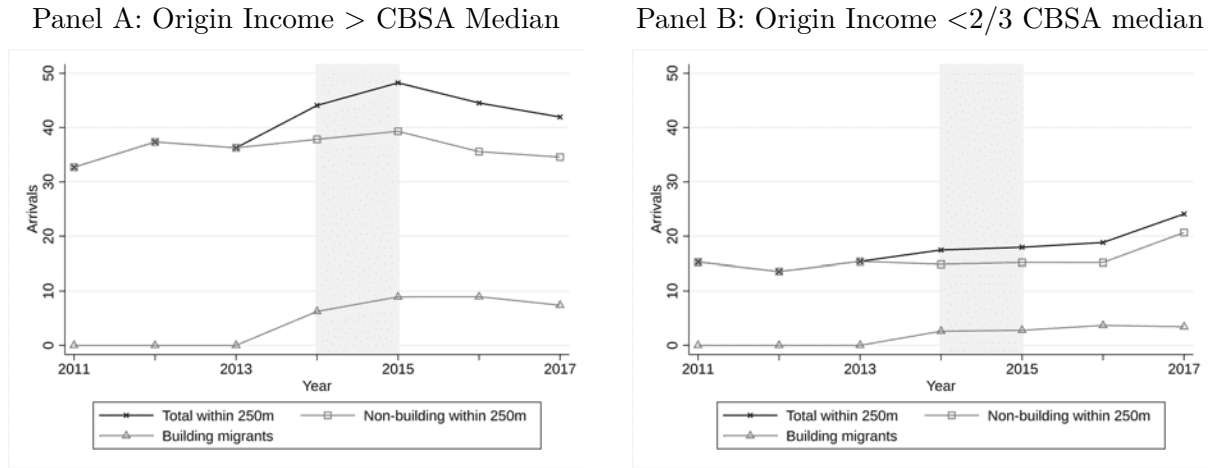
- Autor, D. H., Palmer, C. J., and Pathak, P. A. (2014). Housing market spillovers: Evidence from the end of rent control in Cambridge, Massachusetts. *Journal of Political Economy*, 122(3):661–717.
- Baum-Snow, N. and Marion, J. (2009). The effects of low income housing tax credit developments on neighborhoods. *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.
- 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.
- Busso, M., Gregory, J., and Kline, P. (2013). Assessing the incidence and efficiency of a prominent place based policy. *American Economic Review*, 103(2):897–947.
- 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.
- Curnette, K., Humphries, S., Makar, S., O’Brien, M., and Wacksman, J. (2016). The zillow group report on consumer housing trends.

- Damiano, A. and Frenier, C. (2020). Build baby build?: Housing submarkets and the effects of new construction on existing rents. Working Paper.
- 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. 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.
- Ding, L., Hwang, J., and Divringi, E. (2016). Gentrification and residential mobility in Philadelphia. *Regional Science and Urban Economics*, 61:38–51.
- 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.
- Ganong, P. and Shoag, D. (2017). Why has regional income convergence in the U.S. declined? *Journal of Urban Economics*, 102:76–90.
- Glaeser, E. and Gyourko, J. (2018). The economic implications of housing supply. *Journal of Economic Perspectives*, 32:3–30.
- Glaeser, E. L. and Ward, B. A. (2009). The causes and consequences of land use regulation: Evidence from greater boston. *Journal of Urban Economics*, 65(3):265–278.
- 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.
- 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. (2020). Measuring housing stability with consumer reference data. *Demography*, 57:1323–1344.
- 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.

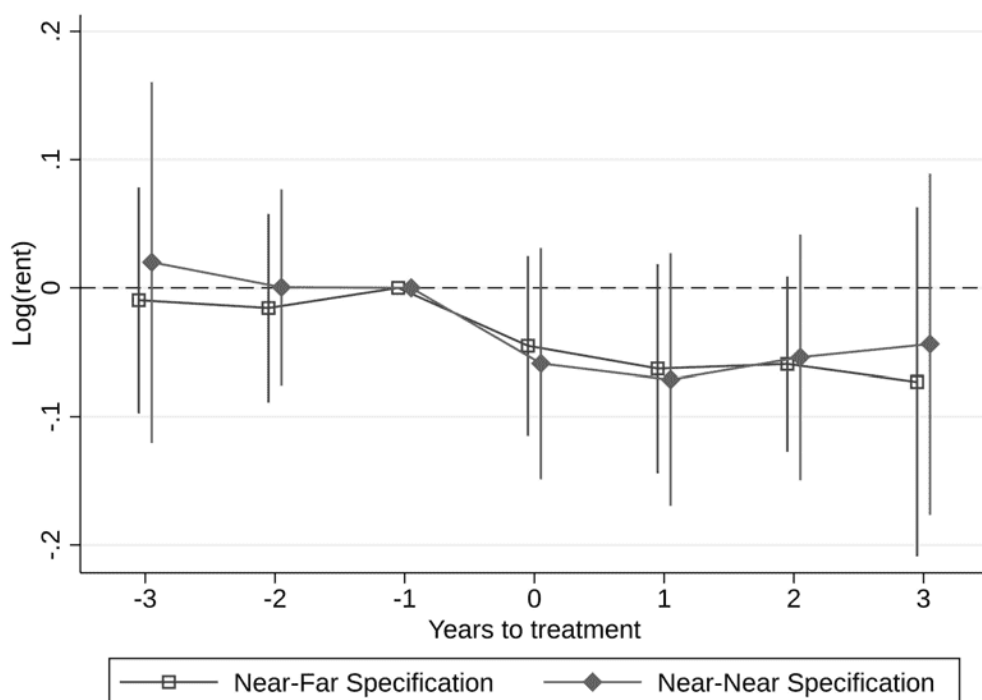
- Rae, A. and Sener, E. (2016). How website users segment a city: The geography of housing search in london. *Cities*, 52:140–147.
- Rosenthal, S. (2014). Are private markets and filtering a viable source of low-income housing? Estimates from a repeat income model. *American Economic Review*, 104(2):687–706.
- 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.
- Shoag, D. and Veuger, S. (2018). Shops and the city: Evidence on local externalities and local government policy from big-box bankruptcies. *Review of Economics and Statistics*, 100(3):440–453.
- Sirull, E. (2018). Here’s what renters say is really most important to them. *Apartment Guide*.
- Su, Y. (2020). 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.
- Zillow Group (2017). Homes selected by renters.

Figure 1: Average In-Migration to Area around New Buildings




Note: This figure shows trends in the average number of in-migrants to the area within 250m of a new building. 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 250m 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 250m 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: Rent Event Studies



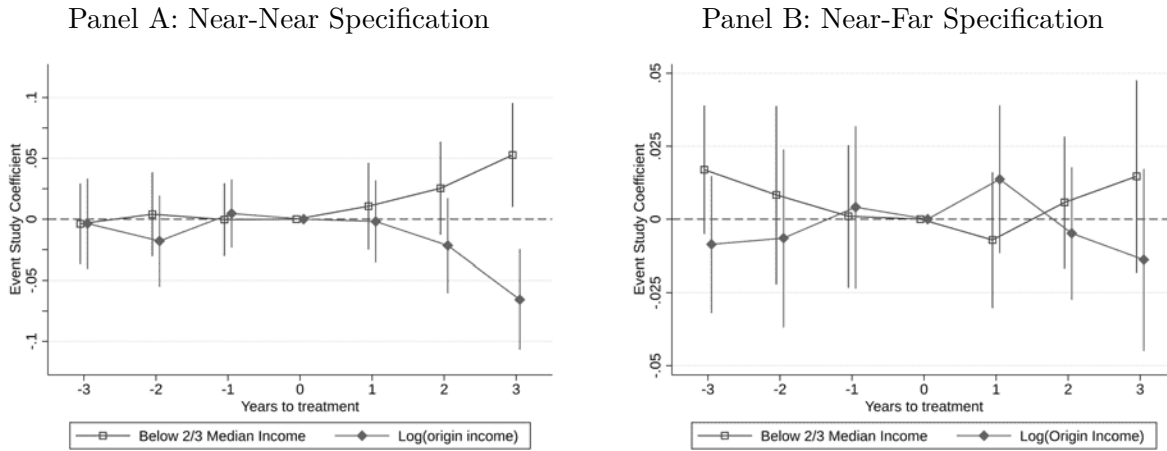
Note: This figure shows the near-far and near-near event studies of the effect of new buildings on nearby rents. In the near-far specification, the treatment group is listings within 250m of a new building, and the control group is listings between 250m and 600m of the same buildings. In the near-near specification, 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 buildings from our final analysis sample. The near-far specification, shown in Equation 1, includes nearest-building \times year and nearest-building \times inner ring fixed effects, as well as controls for bedroom and bathroom counts. The near-near specification, shown in Equation 3, include nearest-building and CBSA \times 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 3: Rent Estimates Across Specifications and Robustness Checks



Note: This figure shows difference-in-differences estimates for the log(rent) outcome across the three empirical strategies, as well as a variety of robustness checks. The thicker bar represents a 90 percent confidence interval, and the thinner, a 95. The parenthetical in each line of the legend indicates the table or appendix table that contains the estimate.

Figure 4: Migration Event Studies



Note: This figure shows event studies of the effect of new buildings on nearby in-migration. The dependent variables are an indicator for whether a migrant originated in a tract below 2/3 of the CBSA median income and the log of the income in a migrant’s origin tract. For the near-near specification, 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). For the near-far, 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 near-near specification includes nearest-building and CBSA \times year fixed effects, and the near-far specification includes nearest-building \times year and nearest-building \times inner ring 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
<u>Building Units</u>				
Units (mean)	198	176	167	165
Units (25th pctile)	85	79	76	68
Units (50th pctile)	162	141	134	117
Units (75th pctile)	287	249	250	256
<u>Nearby Rental Units</u>				
Within 250 m (50th pctile)	349	371	329	320
Within 600 m (50th pctile)	2,010	2,135	1,896	1,843
<u>Rent in Nearby Zillow Listings</u>				
Rent 1 br (mean)	1,784	1,618	1,581	1,546
Rent 1 br (25th pctile)	1,356	1,289	1,272	1,297
Rent 1 br (50th pctile)	1,633	1,543	1,530	1,524
Rent 1 br (75th pctile)	2,045	1,811	1,755	1,740
Rent (mean)	2,129	1,924	1,884	1,867
Rent (25th pctile)	1,667	1,572	1,575	1,605
Rent (50th pctile)	1,942	1,810	1,794	1,836
Rent (75th pctile)	2,441	2,112	2,044	2,061
<u>Buildings by City</u>				
Atlanta	110	31	9	4
Austin	155	100	35	12
Brooklyn	192	104	28	14
Chicago	117	18	5	4
Denver	133	79	25	14
Los Angeles	165	106	28	14
Philadelphia	43	24	4	1
Portland	137	89	27	10
San Francisco	61	27	6	1
Seattle	259	139	51	19
Washington, DC	111	51	13	3
Observations	1,483	768	231	96

Note: Distributions of units in buildings, number of nearby rental units, rents in ZillowTM listings near buildings (within 600 meters), and city locations of buildings for different samples of buildings. The restrictions and definitions used to create the subsamples are described in detail in Section 1.1 of the main text. The number of nearby rental units is estimated based on the tract-level density of rental units in the 2013-2017 ACS.

Table 2: Building Neighborhood Characteristics

	All Incomes		Low-Income			
	No Building	Some Building	No Building	Some Building	2014-6 Building	Final Sample
<u>Household Income</u>						
2000 (\$)	57,879	55,240	46,779	43,872	45,069	42,629
2010 (\$)	58,328	61,513	45,347	45,830	46,669	44,768
2017 (\$)	60,330	76,332	44,897	54,198	54,824	52,783
2000 to 2010 (pct)	-1	11	-4	6	4	5
2010 to 2017 (pct)	2	22	-1	17	16	16
<u>College Degree</u>						
2000 (pct)	25	39	17	29	30	27
2010 (pct)	30	50	21	39	41	37
2017 (pct)	34	60	25	48	50	48
2000 to 2010 (pp)	5	12	4	10	10	10
2010 to 2017 (pp)	4	9	4	9	9	11
<u>Rent</u>						
2000 (\$)	1,001	992	884	869	876	844
2010 (\$)	1,168	1,199	1,044	998	985	956
2017 (\$)	1,271	1,506	1,108	1,233	1,213	1,180
2000 to 2010 (pct)	16	19	17	14	12	13
2010 to 2017 (pct)	7	22	6	21	20	20
Observations	3,218	1,483	2,323	768	231	96

Note: Mean characteristics of census tracts receiving new buildings. 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 (excluding those that are over 25% college students). “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.

Table 3: Baseline Difference-in-Differences Results for Rent Outcome

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

Note: This table shows the effect of new buildings in low-income areas on log rent in nearby units. 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. We include here 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 4: All-Income Difference-in-Differences Results for Rent Outcome

	Near versus far	Near versus near	Triple-difference
After*within 250 (S.E.)	0.016 (0.016)		
After*treated building (S.E.)		-0.017 (0.029)	
After*within 250*treated building (S.E.)			0.019 (0.040)
Treated buildings	103	94	91
Control buildings		31	30
Listing observations	190,020	75,710	201,818

Note: This table shows the effect of new buildings, including all neighborhood income levels, on log rent in nearby units. 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. We include here only new buildings that, apart from the low-income requirement, meet all 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
<i>Panel A: Log(origin income)</i>			
After*within 250 (S.E.)	-0.004 (0.010)		
After*treated building (S.E.)		-0.030 (0.017)	
After*within 250*treated building (S.E.)			-0.027 (0.014)
<i>Panel B: 1(<2/3 median origin)</i>			
After*within 250 (S.E.)	0.002 (0.008)		
After*treated building (S.E.)		0.034 (0.017)	
After*within 250*treated building (S.E.)			0.027 (0.013)
Treated buildings	66	66	66
Control buildings		72	72
Arrival observations	119,400	57,400	186,700

Note: This table shows baseline results for the migration analysis. 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. We include here 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.