Can We Rebuild a City?
Zoning and the Dynamics of Urban Redevelopment*

Vincent Rollet
MIT

February 2025

Abstract

Urban neighborhoods progressively change through the replacement of old buildings with new ones. This paper studies this redevelopment process and how it is influenced by land use regulation. I build a dynamic general equilibrium model of the supply and demand of floorspace in a city, which I estimate using a new parcel-level panel dataset of land use and zoning in New York City and validate using quasi-experimental variation from recent zoning changes. I describe how the city would evolve in the coming decades if zoning restrictions were relaxed. I find that (1) the distortive effects of zoning are concentrated in areas with limited density despite high prices; (2) relaxing zoning constraints primarily benefits households living outside rezoned areas and poorer workers; and (3) zoning constraints are the main determinant of floorspace supply elasticities.

Keywords: City structure, zoning, redevelopment, historical persistence.

*I am immensely grateful to David Atkin, Dave Donaldson, and Tobias Salz for their invaluable advice and constant encouragement, and thank David Autor, Andy Cantu, Pierre-Philippe Combes, Arnaud Costinot, Denise DiPasquale, Edward Glaeser, Vishan Nigam, Nicolas Hommel, Raphaël Lafrogne-Joussier, Lia Petrose, Albert Saiz, Benoit Schmutz-Bloch, Laura Weiwu, Bill Wheaton, participants at the 2024 Conference on Urban and Regional Economics, and seminar participants at MIT for helpful feedback. This work benefited from the generous support of the George and Obie Shultz fund and the Jerry A. Hausman Dissertation Fellowship. I also thank the NYU Furman Center and the Citizens Budget Commission for sharing data. Author contact: vrollet@mit.edu.

1 Introduction

Over time, urban neighborhoods evolve to accommodate a growing population and reallocate land to its most profitable use. While this continuous transformation is key to cities' success, many urban areas struggle to adapt to changing economic conditions. Housing costs have skyrocketed worldwide, some cities failed to transition from productive manufacturing centers to vibrant service-oriented hubs, and the increasing prevalence of working from home threatens previously thriving downtowns.

In urban cores, vacant land is scarce and the built environment mainly changes through redevelopment, the process of tearing down existing buildings and replacing them with new ones. In this paper, I study how cities evolve through this process, using New York City as a case study. In particular, I investigate the extent to which zoning constrains urban growth. Zoning regulations, which limit the size of new buildings and their permitted uses, may be the most important policy lever influencing a city's evolution, and zoning reform is often promoted as an essential measure to tackle the housing crisis.

Studying redevelopment is difficult for three reasons. First, it requires data describing the behavior of developers, the key agents in this process. Which buildings do they demolish, when, and what do they build in their place? This information is typically not readily available. Second, modeling redevelopment decisions is challenging: developers are forward-looking (since redevelopment is expensive and largely irreversible), and construction in one neighborhood affects real estate prices throughout the city. Third, two frictions hinder redevelopment: a policy friction (zoning regulations, which limit developers' choice set) and a technological friction—to construct a new building, a developer must pay a large fixed cost to demolish the existing structure. Separating these two frictions is difficult but essential to understanding the extent to which zoning slows down the evolution of cities.

To overcome these difficulties, I build a parcel-level panel of the built environment in NYC and use it to estimate a dynamic general equilibrium model of the city, predicting how it would evolve in the next decades under different policy scenarios.

My analysis is enabled by newly collected data, presented in Section 2. Using cadastre maps, I divide the city into about 833,000 parcels, which I match to property tax records as well as data from StreetEasy (Zillow's NYC subsidiary). This allows me to track the buildings on each parcel and their characteristics for each year between 2004 and 2022. About 22,000 parcels are redeveloped during my period of interest. For these parcels, I track the redevelopment process using a dataset of building permits and certificates of occupancy. I match this data with the universe of real estate transactions that have taken place in NYC since 2003 and zoning maps for each year.

In Section 3, I highlight several features of the redevelopment process. In NYC, redevelopment is associated with densification. 96% of the buildings that are demolished make way for

taller ones, and new buildings are on average 3.4 times taller than the ones they replace. Conversely, redevelopment seldom happens when it cannot increase the amount of floorspace on a parcel, such as when it is already at the maximum density allowed by zoning. Consistently, analyzing zoning reforms that took place over the past decades, I find that upzoning (relaxing zoning constraints) prompts redevelopment. Finally, redevelopment, in addition to increasing the aggregate supply of floorspace in a city, allows the adjustment of the ratio of residential to commercial floorspace in each neighborhood. Developers will tend to build the type of floorspace that is most valuable, reallocating land to its most profitable use. Zoning regulations that limit allowed uses on each parcel, however, constrain this adjustment.

To estimate the likely effects of broad zoning changes in NYC, I build a dynamic general equilibrium model of the supply and demand of floorspace in a city, described in Section 4. In the model, each parcel of land is associated with a potential developer who can regularly decide whether to redevelop it or not. If the parcel is redeveloped, the developer can decide what to build to replace the existing structure. The new building can be residential or commercial, but its use and size must comply with zoning restrictions. Redevelopment is costly. Some costs are fixed (e.g., those associated with permitting, eviction, and demolition) and will be larger when redeveloping larger structures. Others are variable, and increase with the size of the new building.

The city is populated by workers who choose where to live and work. They consume residential floorspace, and the firms they work for use commercial floorspace in production. Buildings' quality deteriorates over time, making them progressively less valuable, and redevelopment changes the supply of floorspace in each neighborhood. Because of agglomeration forces, amenity and productivity levels in each neighborhood adjust following construction. Floorspace rents adjust so that markets clear. Developers, aware of these dynamics, make forward-looking decisions.

Section 5 describes how I estimate this model. For the demand side, which determines how prices react to changes in floorspace supply, I build on the quantitative spatial economics literature (Ahlfeldt, Redding, Sturm and Wolf, 2015; Redding and Rossi-Hansberg, 2017; Tsivanidis, 2019). To estimate agglomeration forces, I leverage the construction of large new buildings. Comparing areas that witnessed construction earlier vs. later, I estimate how local rents and floorspace quantities respond to such shocks. I then map the implied elasticities of rents to floorspace to spillover parameters. The agglomeration elasticities I measure are comparable to available estimates in the literature, and I do not find evidence of strong spillovers across land uses (i.e., of residential buildings on firms or commercial buildings on residents).

For the supply side, which models the behavior of developers, I rely on methods developed in empirical industrial organization and their application to urban economics (Rust, 1987; Murphy, 2018; Almagro and Domínguez-Iino, 2024; Hsiao, 2024). Specifically, I proceed in three steps. First, using real estate sales data, I use hedonic regressions to estimate the value of buildings as

a function of their characteristics. This enables me to estimate the present value of properties in NYC and assess how redevelopment could increase their worth. Then, I estimate variable costs using a revealed preferences approach. When redevelopment takes place, I observe the choices made by developers as well as the floorspace prices and zoning regulations they faced. This provides information about the cost schedule faced by developers. Finally, I recover fixed costs through an extension of Rust (1987)'s nested fixed point approach, leveraging the relationship between the probability that a parcel is redeveloped during a given year and the profit that developers could expect if they redeveloped the parcel. I find that fixed costs increase sharply with the size of the building to demolish, and are also higher in denser neighborhoods.

I show that the estimated model can adequately predict the probability that a parcel will be redeveloped in a given year and what developers build if they proceed with redevelopment. In particular, the model matches well with my quasi-experimental estimates of the effect of recent upzonings.

The estimated model allows me to predict NYC's evolution under a wide range of zoning scenarios. In Section 6, I present results from several counterfactual exercises. I first evaluate how zoning regulations constrain the city's growth. In the absence of zoning changes, I estimate that the supply of floorspace in NYC would grow by 17% between 2019 and 2060. If zoning regulations were completely removed (excluding parcels that are landmarked, in historic districts, or in flood zones), NYC's growth rate would more than triple and it would have 71% more floorspace by 2060. An ambitious but realistic upzoning of NYC around its transit stations could deliver about a sixth of the floorspace supply boost that removing zoning would achieve.

The effects of relaxing zoning are highly heterogeneous. Relaxing zoning has very limited effects in neighborhoods where floorspace is inexpensive. There, redevelopment is usually unprofitable and therefore rare, and new buildings tend to be small, often below zoning limits. Upzoning parcels that are already densely built is also largely ineffective. Even in the absence of zoning constraints, their redevelopment is unlikely as it would entail particularly high demolition costs. The neighborhoods whose growth is most distorted by zoning are those with limited density despite high floorspace prices. Accordingly, I find that relaxing zoning limits in Western Brooklyn and Northern Queens would lead to large increases in floorspace supply, but upzoning most parts of the Bronx or Midtown Manhattan would not substantially boost construction.

A simple static framework that ignores the large fixed costs associated with redevelopment would vastly overstate the effects of relaxing zoning, predicting that removing zoning would lead to more than a doubling of the city's size. While zoning imposes a tight constraint on urban growth, cities' ability to adjust their structure is also strongly limited by technological frictions to redevelopment.

My estimates suggest that relaxing zoning would only lead to moderate rent decreases in NYC (-17% by 2060 if zoning were completely removed), as additional construction would trigger migration to the city and increase floorspace demand there. This limits the welfare gains of

incumbent New Yorkers, and a large share of the gains of relaxing zoning in NYC would accrue to workers in the rest of the country. I further find that poorer workers would disproportionally benefit from zoning reform.

Finally, I estimate floorspace supply elasticities at the neighborhood level and for different time horizons. I find that under its current zoning regulations, NYC has a very low aggregate floorspace supply elasticity, reaching 0.18 at a 40-year horizon. Supply elasticities vary widely across neighborhoods and are essentially determined by zoning stringency. If zoning regulations were removed, the city's floorspace elasticity would more than triple, and local supply elasticities would be mostly determined by neighborhoods' density.

Contribution to the literature. Understanding the economic forces underpinning the allocation of land uses in cities has long been a core question of urban economics. Early work by Burgess (1925) and Hoyt (1939) highlighted how cities tend to allocate land to their most profitable use. This insight was systematized in monocentric city models, following the seminal work of Alonso (1964), Mills (1967), and Muth (1969). In particular, Ogawa and Fujita (1980) and Lucas and Rossi-Hansberg (2002) describe the equilibrium distribution of residential vs. commercial space in a symmetric city. Quantitative spatial models, pioneered by Allen and Arkolakis (2014) and Ahlfeldt et al. (2015), enable the analysis of city structure with more realistic geographies. Allen, Arkolakis and Li (2015) study the allocation of land uses in this framework and Heblich, Redding and Sturm (2020) show that quantitative spatial models can satisfyingly rationalize the growth and reorganization of cities following the development of railways.

Such studies of city structure typically focus on long-term, steady-state equilibria and assume that land uses in cities can be adjusted frictionlessly. However, cities change very slowly and their evolution is severely constrained, especially as they mature. It is prohibitively expensive to add more floorspace to a given parcel if there is already a skyscraper there, and zoning may render some land use changes fully illegal. Hence, I focus in this paper on the transition dynamics of cities and study how zoning affects them, which is more informative for urban policy. This relates my study to Hsiao (2024), who builds a dynamic model of developers to study coastal development in Jakarta, Gechter and Tsivanidis (2023), who study the redevelopment of industrial land in Mumbai, Peng (2023), who evaluates long-term impacts of recent rezonings in NYC, and Greaney, Parkhomenko and Van Nieuwerburgh (2025), who build a rich dynamic model to study how metropolitan areas progressively respond to shocks. These studies use neighborhoods as the unit of analysis and conceptualize redevelopment as an increase in a neighborhood's capital stock. In contrast, I study redevelopment at the parcel level. This allows me to disentangle the two primary obstacles to redevelopment—demolition costs and zoning restrictions—both of which vary from one parcel to the next. Leveraging microdata on developers' behavior allows for a finer analysis of the production function for floorspace. This connects my study to Soltas (2024), who uses similarly granular data to study the effect of subsidies on

low-income housing construction.

By evaluating the role of zoning in shaping city structure, I contribute to an active literature assessing the impact of land use regulations on cities. This literature has shown that zoning reduces housing supply, raises prices, and results in a costly misallocation of workers across space (Glaeser, Gyourko and Saks, 2005; Gyourko and Molloy, 2015; Ganong and Shoag, 2017; Glaeser and Gyourko, 2018; Hsieh and Moretti, 2019; Ospital, 2022). Several studies have also leveraged recent zoning changes to show that relaxing land use regulation leads to additional construction (Anagol, Ferreira and Rexer, 2024; Büchler and Lutz, 2021; Liao, 2022; Peng, 2023). Observed rezonings, however, are typically limited in scope and are likely enacted in areas where planners believe they will be effective. By studying developers' behavior at the micro level in a dynamic framework, I contribute to such evaluations of zoning reform by allowing extrapolation beyond these case studies. Indeed, the model developed in this paper can be used to predict the effect of a wide range of zoning reforms over long time horizons. I find that upzoning only leads to substantial construction in some neighborhoods, and the effect of such policy changes takes decades to materialize.

Furthermore, the existing work on zoning has overwhelmingly focused on housing and density restrictions. I extend my analysis to commercial real estate and use regulations (which restrict where commercial and residential buildings can be constructed). This extended scope is valuable for multiple reasons: commercial buildings represent a large share of the total floorspace stock (35% in NYC), their location is a key determinant of job access and commuting times, and the separation of residential areas from commercial ones on zoning maps has strongly influenced the structure of American cities (Hirt, 2015). This broader focus relates my work to Kulka, Sood and Chiumenti (2024), who also investigate the multifaceted nature of zoning by studying the interaction between different types of regulations limiting housing density.

Finally, I contribute to a literature measuring housing supply elasticities in cities (Saiz, 2010; Murphy, 2018; Combes, Duranton and Gobillon, 2021; Baum-Snow and Han, 2024) by providing granular estimates of floorspace supply elasticities at different time horizons and quantifying the extent to which they are distorted by zoning.

2 Data and Context

With over eight million inhabitants, NYC is the largest city in the United States. The total value of its real estate exceeds three trillion dollars, and it has been slowly growing over the past decades. To analyze its evolution, I assembled data from a wide range of sources. This section describes the key datasets I rely on, and Appendix A provides additional details.

A parcel-level panel of NYC buildings. I track the evolution of land use in NYC at the smallest possible scale: the parcel. Such granular data is valuable as developers make decisions at this

level, and zoning regulations are parcel-specific. However, parcel-level panels of land use are not readily available. The closest existing datasets consist of repeated cross-sections of all buildings in a given city. Redevelopment causes some buildings to disappear from the cross-section and others to appear. To construct a panel dataset of land use, one must link new buildings with the old ones that were demolished to make way for them.

To do so, I rely on past and current cadastre maps of the city. They allow me to divide NYC into 833,000 parcels and find the property tax identifiers of the buildings located on them for each year between 2004 and 2022. Using property tax records, as well as data from StreetEasy, the NYC subsidiary of Zillow, I then associate these buildings with a wide range of characteristics.

Over my period of study, about 22,000 parcels have been redeveloped. To understand when redevelopment took place, I match my panel with a dataset of building permits (indicating when redevelopment started) and a dataset of certificates of occupancy (documents issued by the city certifying that a new building is fit for use).[1] Construction projects are usually completed quickly, with certificates of occupancy typically granted two years after the corresponding building permit (see Appendix Figure C.2).

Real estate prices and rents. To measure floorspace prices throughout the city and over time, I rely on the Property Sales Files, a dataset provided by the NYC Department of Finance that documents the universe of real estate transactions that took place in NYC between 2003 and 2022.

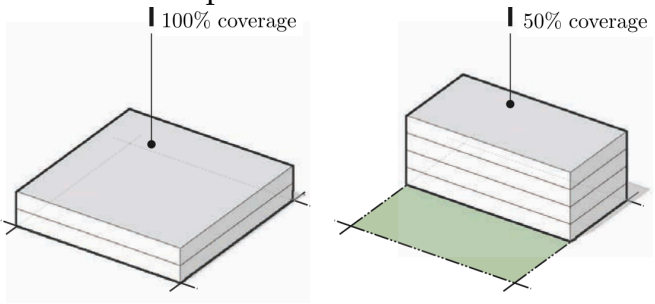
For a subset of buildings, I can further recover yearly rent estimates from the annual Notices of Property Values (NOPVs) sent to property owners to notify them of their property tax obligations. These documents are used in Li (2022) to build building-level rent measures covering 2003–2013. By scraping NOPVs issued in recent years, I extend Li's dataset up to 2021.

Zoning constraints. Finally, I extract from NYC's zoning resolution the regulatory constraints imposed on different parcels. The zoning map, of which an extract is presented in Appendix Figure B.1, associates each land parcel with a zoning district, indicating the allowed land uses on the lot, as well as the maximum allowed bulk of new buildings. Density in NYC is mostly regulated through Floor Area Ratio (FAR) limits. The FAR of a parcel is defined as the floorspace area of buildings on that parcel divided by the parcel's land area. Figure 1(a) shows two parcels with a FAR of 2.0, and Figure 1(b) shows the maximum FAR limits associated with four of the city's zoning districts. For instance, in R1 districts, which are designed for single-family homes, only residential buildings can be constructed, and their FAR must not exceed 0.5.

[1] I consider that a parcel has been redeveloped when a new building is constructed. I also tracked conversions (usually of commercial buildings to residential ones), which account for a small share of total redevelopment activity (see Appendix Figure C.1) and are therefore ignored in the quantitative exploration below. Renovations and building extensions are also excluded from the scope of this study as the amount of floorspace added to the city through them is an order of magnitude lower than that added through new construction.

Figure 1: Illustration of zoning restrictions in NYC

(a) Two parcels with a FAR of 2.0



(b) Examples of zoning districts in NYC

Zone	Max res. FAR	Max com. FAR
R1	0.5	-
R5	1.25	-
M1-4	-	2
C6-6	10	15

Notes: Panel (a) shows an illustration of the measurement of the Floor Area Ratio (FAR) by the NYC Department of City Planning. It is defined as "the ratio of total building floor area to the area of its zoning lot." Panel (b) shows the maximum FAR limits for residential and commercial buildings in four zoning districts. R1 is a residential district for single-family homes. R5 is a medium-density residential district. M1-4 is a medium-density light manufacturing district that allows a broad range of commercial uses, including offices, but not residential uses. C6-6 is a high-density district that allows both residential and commercial uses but favors the latter.

NYC's current zoning resolution was initially adopted in 1961 and largely reflects the city's organization in the early 1960s. Commercial zoning districts are concentrated in the high-density neighborhoods of the lower half of Manhattan and Downtown Brooklyn, while more peripheral neighborhoods are almost exclusively zoned for residential purposes with strictly restricted densities (see Appendix Figures B.2 and B.3).

The average parcel in NYC has a FAR of 1.1 and a maximum allowed FAR of 1.4 (see Appendix Figure C.3). A third of the parcels have a FAR over the maximum, often because they were built before 1961 and grandfathered in. About a quarter have FARs within 0.25 points of the limit, making density limits binding or near binding on most of the city's parcels.

While NYC's zoning code is now considered very restrictive, this was not the view or desire of its designers. Urban planners of the mid-20th century believed the city had nearly reached its ultimate size. Their main aim was to guide NYC's limited future growth, particularly by promoting tower-in-the-park architecture and further separating commercial development from residential areas.

Contemporary planners hold markedly different views on effective urban planning, typically favoring mixed-use developments and opposing tower-in-the-park projects. Yet, despite these radical shifts in planning principles, the numerous amendments to NYC's zoning resolution since 1961 have mainly brought minor changes, and the city's current zoning map almost perfectly aligns with its 1960s version (see Appendix Figure B.4). In recent decades, however, several neighborhoods have been upzoned, meaning that FAR restrictions have been substantially relaxed there. These areas offer valuable case studies to investigate the effects of zoning reform.

In Appendix Section B.1, I provide a brief history of zoning in NYC. A striking regularity is that the issues prompting planners to implement various zoning regulations often disappear in subsequent decades. Yet, these regulations tend to persist far longer than planners initially

anticipate or intend.

3 Motivating Facts

The granular datasets I collected allow me to describe developers' behavior at the micro level. I highlight the data's most striking features in this section and provide additional descriptive statistics in Appendix C.

Redevelopment is associated with densification. Redevelopment entails high fixed costs. Indeed, to construct a new building, one must demolish existing structures, forgo rent payments during construction, apply for building permits, wait for existing tenants to leave, etc. Hence, redevelopment is only financially viable when new structures are significantly more valuable than the buildings they replace. A developer can increase the value of a parcel by increasing the quality of the building on it. Indeed, I find that older structures are much more likely to be redeveloped than newer ones. However, new buildings are more valuable than old ones mostly because they are taller: Figure 2(a) compares the FAR of old structures with those they replaced. 96% of redevelopment events increased the amount of floorspace on the redeveloped parcel, and on average, new buildings had a 3.4 times higher FAR than those they replaced.

Figure 2: Redevelopment happens when it can substantially increase a parcel's floorspace



Notes: Panel (a) compares, for redeveloped parcels, the FAR of new structures with the FAR of the structures they replaced. The dashed line corresponds to the 45-degree line. Panel (b) shows the probability of observing a redevelopment event in the panel for different bins of the distance between the FAR of the existing structure and the maximum FAR allowed by zoning.

A corollary pattern is that redevelopment tends to only take place when it is possible to replace an existing structure with a larger one. Figure 2(b) shows that buildings that are at or above the maximum allowed FAR are very rarely redeveloped. Indeed, it is unlikely that redevelopment would increase the value of these structures enough to cover the fixed costs of redevelopment. On the other hand, buildings far from the maximum allowed FAR are much more likely to be redeveloped.

Relaxing zoning prompts redevelopment. The extent to which zoning distorts redevelopment can be directly evaluated by analyzing recently rezoned parcels. Rezonings are politically fraught, require approval from a wide range of stakeholders, and can take decades. Therefore, when the Department of City Planning initiates the process of rezoning a neighborhood, there is considerable uncertainty regarding which parcels will be rezoned, how their FAR and use limits will change, and when those changes will take effect.[2]

While the parcels that are upzoned are typically ripe for redevelopment and represent a selected sample, one can reasonably consider the exact timing of rezonings to be exogenous. Therefore, to estimate the effect of upzoning on redevelopment, I isolate parcels within areas rezoned between 2004 and 2022 that were upzoned, meaning they experienced an increase in their maximum allowed FAR by more than 0.25 FAR points.[3] I then compare parcels that were upzoned earlier vs. later using the procedure of De Chaisemartin and d'Haultfoeuille (2020). Figure 3 presents event studies estimates of the effects of upzoning on the built FAR.

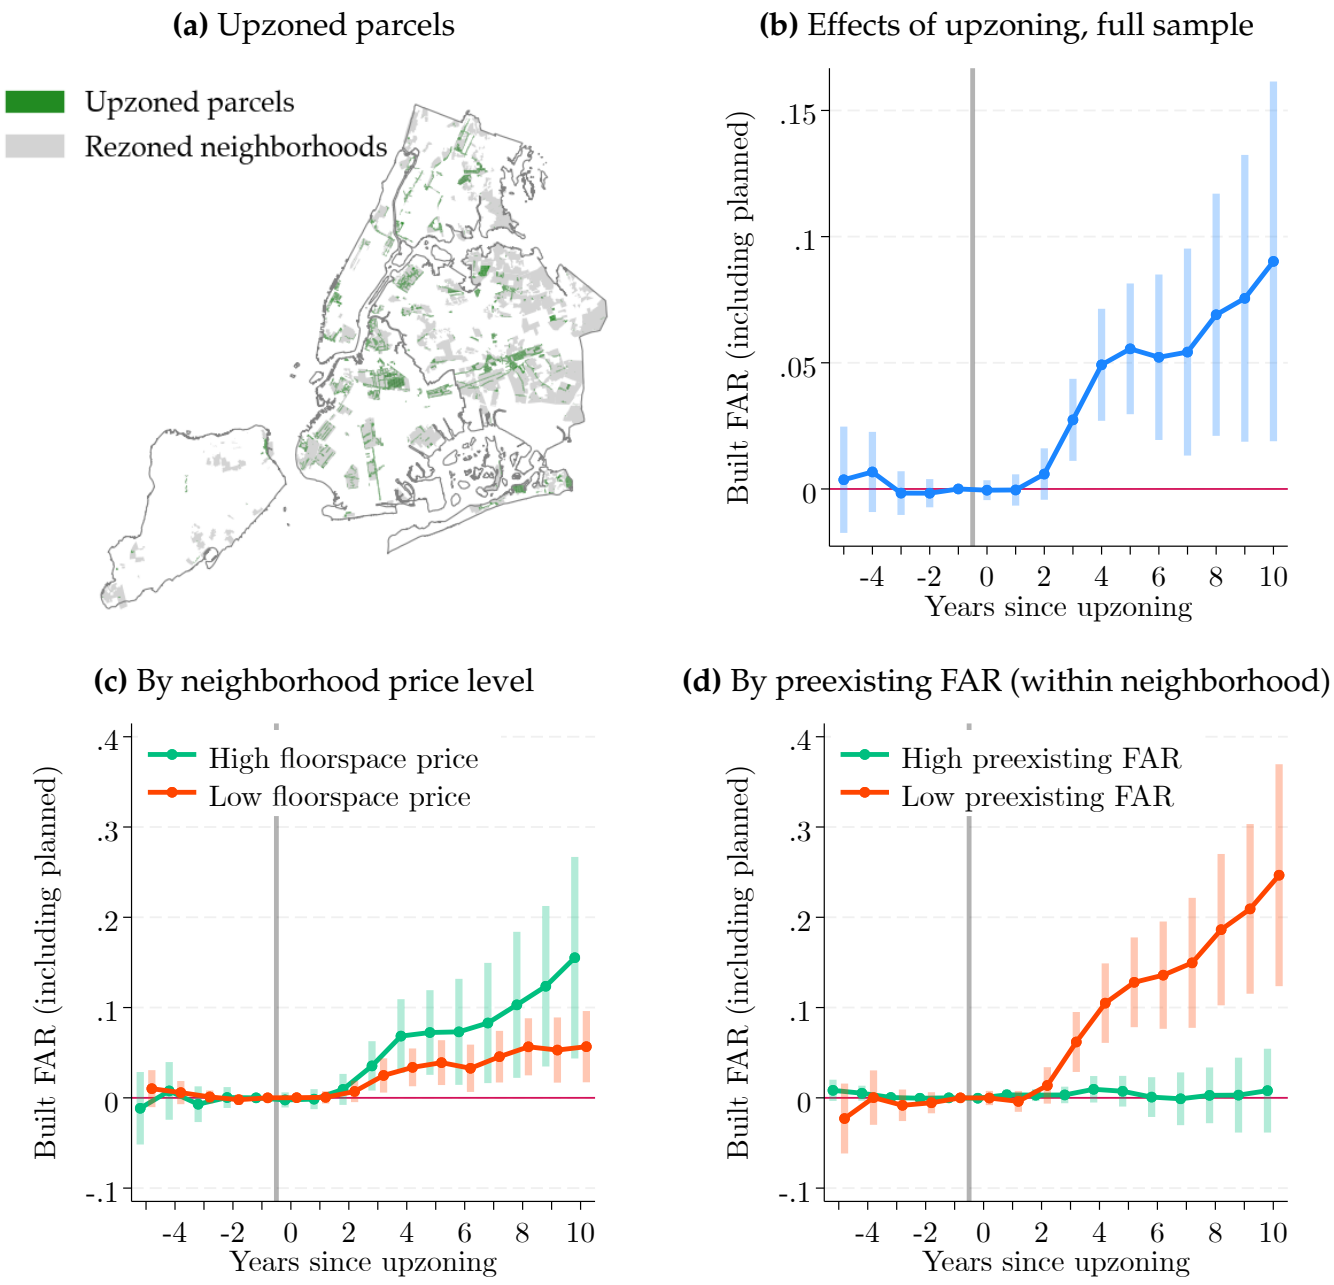
Relaxing zoning limits leads to increases in development activity and floorspace supply. An increase in the maximum allowed FAR by 1.1 units is associated with an increase in the average FAR of 0.9 units over a 10-year horizon (panel b). This represents a sizeable increase—in the year before the rezoning, upzoned parcels had an average FAR of 1.3.

Therefore, zoning appears to be a strong constraint on floorspace supply in the city. This result aligns with previous studies of upzonings in NYC (Liao, 2022; Peng, 2023), and is consistent with the wide gap between the price of floorspace and its marginal construction cost. In 2019, reported construction costs in NYC ranged between $160 and $600 per sq. ft of floorspace (see, e.g., RSMeans, 2020; New York Building Congress, 2019). At the same time, the price of floorspace in NYC averaged around $700 per sq. ft, reaching much higher levels in some neighborhoods of Manhattan. This large wedge, noted by Glaeser, Gyourko and Saks (2005), would be arbi-

[2]For instance, discussions to rezone the Gowanus Canal area began in the mid-2000s during the Bloomberg administration. Local community groups opposed rezoning proposals, fearing that the changes would lead to gentrification and worsen the canal's flooding and contamination issues. The neighborhood's rezoning plan was finally approved in late 2021, after nearly twenty years of negotiation.

[3]In the analysis sample for the event study, I exclude parcels that were upzoned as part of small rezonings, defined as rezonings affecting fewer than 30 parcels. These small rezonings are often requested by developers and may be processed more quickly than larger ones, making their occurrence likely endogenous. Large rezonings, conversely, tend to be initiated by city planners and follow a longer approval process. Approximately 10% of the redevelopment events occur on parcels upzoned as part of a large rezoning, and 1% on parcels upzoned as part of a small rezoning. Parcels in the estimation sample are mapped in Figure 3(a).

Figure 3: Effects of upzoning on built FAR



Notes: This figure presents estimates of the effects of upzoning on built FAR. Panel (a) shows (in gray) the areas that were rezoned between 2004 and 2022, excluding small rezonings affecting fewer than 30 parcels. The estimation sample comprises parcels within these areas that were upzoned (i.e., that saw their maximum allowed FAR increase by more than 0.25 points), plotted in green. Panel (b) shows the effect of upzoning on the built FAR (including the planned floorspace of buildings under construction) for the full sample. In Panel (c), I split the estimation sample between neighborhoods with floorspace prices above/below the median. In Panel (d), I split the sample between parcels whose built FAR in 2004 was above/below the median of their neighborhood. Treatment effects are computed using the procedure of De Chaisemartin and d'Haultfoeuille (2020), comparing parcels upzoned earlier vs. later. Observations in the regression are weighted by parcel size, and standard errors are clustered at the city block level. The average parcel in the sample was upzoned by 1.1 FAR points.

traged by developers if they could easily adjust the floorspace supply of the city. Nonetheless, the results of Figure 3(b) show that when zoning restrictions are relaxed, newly built floorspace

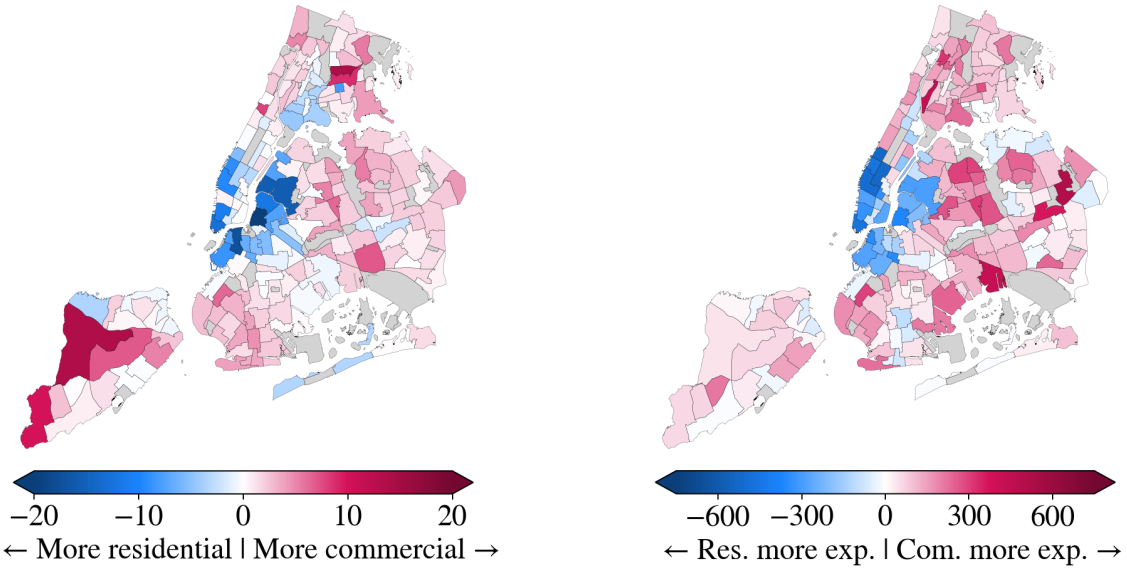
represents only a fraction of the newly allowed floorspace. This is consistent with the need to pay large fixed costs to redevelop a parcel and expand its FAR.

The effects of upzoning are strikingly heterogeneous. Estimating the effects of upzonings separately for neighborhoods with high vs. low floorspace prices (panel c) reveals that the effects of upzoning on construction are almost entirely driven by upzonings in expensive neighborhoods. Upzoning parcels in areas where floorspace prices are low yields very little new construction—there, the reward for increasing floorspace supply is usually too low to cover redevelopment costs. Furthermore, the effects of upzoning are fully concentrated on parcels that had a relatively low FAR before upzoning (panel d). Upzoning parcels that already have tall buildings typically does not result in additional construction, due to the high fixed costs associated with redeveloping these structures.

Redevelopment reallocates land to its most profitable use. Land uses in NYC are becoming increasingly mixed. Figure 4(a) shows that over the past decades, the (mostly commercial) central neighborhoods of the city have shifted to become more residential, while its (mostly residential) peripheral neighborhoods have witnessed increases in their share of commercial floorspace.

Figure 4: Redevelopment reallocates land to its most profitable use

(a) Change in commercial share, 2004-2022 **(b)** Commercial-residential price differential



Notes: Panel (a) shows the change in the share of commercial floorspace in total floorspace between 2004 and 2022. Panel (b) shows the price difference between a sq. ft of commercial floorspace (office or retail) and a sq. ft of residential floorspace. Data is aggregated at the neighborhood tabulation area level.

These city-level trends are consistent with the financial incentives that developers face. Indeed, in NYC's central neighborhoods, residential floorspace is on average 30% more expensive than commercial floorspace, and in peripheral neighborhoods, commercial floorspace is on av-

erage 30% more expensive than residential floorspace—see Figure 4(b). Overall, redevelopment therefore reallocates land to its most profitable use.

These trends are consistent with NYC's shift from a manufacturing hub to a service-oriented metropolis: since 1961, the share of manufacturing in employment in New York has been divided by 10, and the central areas of the city have become less noisy, less polluted, and generally more attractive for residents. Between 1950 and 2019, rents in the central areas of the city grew 2.5 times more than in its more residential neighborhoods (see Appendix Figures B.5 and B.6). Businesses, on the other hand, have been migrating out of the central business district.

Smartphone location data, analyzed in Appendix Figure B.7, suggests that this trend will accelerate with the rise of working from home. In the aftermath of the 2020 pandemic, visits to commercial locations in the center of NYC have decreased relative to those in peripheral locations, consistent with workers spending more time in residential neighborhoods.

Existing zoning regulations, however, slow down the reallocation of land uses. Indeed, zoning favors the construction of commercial buildings in the center of the city and usually disallows such construction in peripheral neighborhoods (see Appendix Figure B.8). The bindingness of zoning rules limiting parcels' residential or commercial use can be directly tested using event studies analogous to that of Figure 3. In Appendix Figure F.1, I show that rezonings that newly allow a parcel to have a residential (or commercial) use lead to an increase in the probability that the parcel has such a use.

4 A Model of Urban Evolution

To better understand the city's slow evolution through redevelopment and evaluate the extent to which zoning distorts this process, I build a dynamic general equilibrium model of the city. In the model, new floorspace is supplied by forward-looking developers whose decisions depend on prices, zoning regulations, and the characteristics of existing buildings. Workers consume residential floorspace, while firms use commercial floorspace in production.

4.1 Supply

The model's supply side is designed to capture three key aspects of the redevelopment process. First, replacing old buildings with new ones entails large fixed costs, which can vary with the characteristics of the building to demolish. Second, what developers build is restricted by zoning and influenced by market prices. Third, as for other forms of investment, redevelopment decisions are made dynamically. Buildings deteriorate slowly over time, and the price of floorspace evolves from year to year as a function of future expected rents, leading developers to optimize when they tear down an existing structure to replace it with a new one.

Setup. The city comprises a fixed set of parcels, each associated with a potential developer. Each year, developers decide whether to redevelop their parcel or not, taking an action a in $\mathcal{A}_{it}^{\text{allowed}} \subset \{\emptyset, R, C\}$. $\emptyset$ corresponds to not redeveloping the parcel, and R (resp., C) corresponds to building a new residential (resp., commercial) structure on it. If a developer decides to build, they choose h^{new}, the FAR of the new structure.

The set of available actions $\mathcal{A}_{it}^{\text{allowed}}$ is limited in two ways. First, to redevelop a parcel, a developer has to buy it. Redevelopment, therefore, can only happen when the parcel is transacted on the real estate market, which happens with probability $p_i^{\text{transaction}}$ each year.[4] Second, new construction has to comply with zoning. If the parcel is zoned for residential uses only, the developer cannot decide to build a new commercial building on it.

Developers make decisions such as to maximize their discounted sum of profits,

$$\max_{a_{it} \in \mathcal{A}_{it}^{\text{allowed}}} \mathbb{E}_t \left[\sum_{t'=t}^{\infty} \left(\frac{1}{1+r} \right)^{t'} (\pi_{it'}^{a} + \epsilon_{it'}^{a}) \times K_i \right], \tag{1}$$

Where π_{it}^{a} corresponds to the payoff of taking action a on parcel i in year t, and ϵ_{it}^{a} is a parcel-period-action-specific profit shock. To facilitate interpretation, these payoffs and shocks are measured per sq. ft of land. Hence, the total payoff of choosing a_{it} is $(\pi_{it'}^{a} + \epsilon_{it'}^{a}) \times K_i$, where K_i is the land area of parcel i.

Flow payoffs are all measured relative to the status quo (choosing not to redevelop), and are given by equation (2),

$$\pi_{it}^{\emptyset} = 0, \qquad \pi_{it}^{\theta} = \max_{h^{\text{new}} \leq \bar{h}_{\theta it}} \left[\underbrace{P_{i,t}(h^{\text{new}}, \mathbf{x}_{it}^{\text{new}}) - P_{it}(h_{it}^{\text{old}}, \mathbf{x}_{it}^{\text{old}})}_{\text{Change in property value}} - \underbrace{[VC_{it}^{\theta}(h^{\text{new}}) + FC_{it}]}_{\text{Cost of redevelopment}} \right], \tag{2}$$

Where $\theta \in \{R, C\}$. When redevelopment occurs, developers choose the FAR of the new building h^{new} such as to maximize redevelopment profits subject to the FAR of the new building being below the zoning limit $\bar{h}_{\theta it}$. The developer's profit corresponds to the increase in the parcel's value brought by redevelopment minus redevelopment costs, which can be fixed or variable. $\mathbf{x}_{it}^{\text{new}}$ (resp., $\mathbf{x}_{it}^{\text{old}}$) corresponds to the characteristics of the new (resp., old) building, including, for instance, its age, type, and condition. $P(h, \mathbf{x})$ corresponds to the discounted sum of future rents derived from a building—as time passes, buildings age and their value $P_{it}(h_{it}^{\text{old}}, \mathbf{x}_{it}^{\text{old}})$ decreases, making redevelopment increasingly profitable. Buildings that are shorter and with less valuable

[4] Buildings in NYC are usually not directly redeveloped by their owners. Rather, the redevelopment process tends to start with a developer purchasing a building and applying for a building permit. Appendix Figure D.2 shows that the probability of a parcel being sold increases in the years leading to a redevelopment event and reaches about 40% in the year before the issuance of a building permit, a figure much higher than the unconditional 5.6% probability of a parcel being sold in any given year. The probability of a redeveloped parcel being sold in the year redevelopment starts or in the two preceding years is 72.5%.

characteristics will be more profitable to redevelop and hence are more likely to be demolished.

Variable costs correspond to construction costs that increase with the size of the new building h^{new}. Fixed costs are those associated with demolition, vacating existing tenants, and permitting. They tend to increase with the size of the building to demolish and the density of the neighborhood in which it is located.

I assume that developers have perfect foresight. Furthermore, I rule out strategic interactions between them. This is justified by the very large number of parcels in the city, as well as the competitive nature of the real estate development industry in NYC, which encompasses thousands of firms and has a low degree of concentration (see Glaeser et al., 2005).

Parametrization. To estimate the model, I parametrize price and cost functions. For prices, I assume that the expected value of a parcel (per sq. ft of land) is a function of its geographic location and the characteristics of the building occupying it, with

$$P_{it}(h_{it}, \mathbf{x}_{it}) = \underbrace{\rho^{\theta}_{nt}}_{\text{Location premium}} \times \underbrace{Q(\mathbf{x}_{it})}_{\text{Building quality}} \times \underbrace{h_{it}}_{\text{Building FAR}}. \tag{3}$$

ρ^{θ}_{nt} captures the price of floorspace of type θ (residential or commercial) in neighborhood n in year t, and $Q(\mathbf{x}_{it})$ reflects the building's quality, which is a function of its characteristics, including its FAR h_{it}.

Studies estimating the production function for housing have found that it can be satisfyingly approximated by a Cobb-Douglas aggregating land and capital for both small and tall structures (Epple, Gordon and Sieg, 2010; Combes, Duranton and Gobillon, 2021; Ahlfeldt and McMillen, 2018). Hence, I assume that that investing k units of capital on a sq. ft of land allows to build a structure of FAR $h = k^{\zeta}$. The cost of a unit of capital is given by $\alpha^{\theta}_{it} e^{\eta_{it}}$, where η_{it} is a construction cost shock. Hence, the variable cost function is given by

$$VC^{\theta}_{it}(h^{\text{new}}_{it}) = \alpha^{\theta}_{it}(h^{\text{new}}_{it})^{1/\zeta} e^{\eta_{it}}. \tag{4}$$

Some fixed costs of redevelopment are independent of the height of the preexisting building, while others increase with the amount of floorspace to demolish. Furthermore, demolition costs are likely higher for protected parcels (either as landmarks or because they are in historic districts) and in denser neighborhoods. Hence, I parametrize fixed costs as:

$$FC_{it} = (\delta^{0} + \delta^{\text{demolition}} h^{\text{old}}_{it} + \delta^{\text{density}} \bar{h}_{n(i)t}) \delta^{\text{protected}} \tag{5}$$

δ^{0} reflects the fixed cost (per sq. ft of land) associated with starting the construction of a building on vacant land that is unprotected (i.e., not landmarked or in a historical district). $\delta^{\text{demolition}}$ measures how fixed costs grow with the FAR of the structure to demolish, δ^{density} how they

grow with the average FAR of parcels in the same neighborhood, $\bar{h}_{n(i)t}$, and $\delta^{\text{protected}}$ quantifies the extent to which protected parcels are harder to redevelop.

Additional assumptions. To facilitate estimation, I make three additional assumptions. First, I assume that redevelopment is a terminal decision, i.e., that once a building has been redeveloped, it will not be redeveloped again in the future. This assumption simplifies estimation as it turns the developer's dynamic decision process into an optimal stopping problem. Because the median age of redeveloped buildings is 85 years, and I only simulate the city's evolution over a 40-year period in counterfactuals, this assumption is unlikely to substantially affect my results.

Second, I assume that the profit shocks ϵ^a_{it} are distributed i.i.d. according to a Type 1 Extreme Value distribution with a mean of zero and a scale parameter $\sigma_{it} = \sigma^0_\epsilon + \sigma^{\text{demolition}}_\epsilon h^{\text{old}}_{it} + \sigma^{\text{density}}_\epsilon \bar{h}_{n(i)t}$. $\sigma^{\text{demolition}}_\epsilon$ and $\sigma^{\text{density}}_\epsilon$ are both expected to be positive, capturing the fact that the shocks facilitating or hindering redevelopment are more variable for parcels with larger buildings or in dense neighborhoods.

Finally, I assume that each period, developers choose an action in $\mathcal{A}^{\text{allowed}}$ after observing profit shocks ϵ^a_{it}, but they only observe the cost shock η_{it} if they make the decision to redevelop. After observing η_{it}, the developer chooses the amount of floorspace they want to provide. This allows me to estimate the model in steps, as described in Section 5.

4.2 Demand

Developers' choices lead to adjustments in the supply of floorspace in each neighborhood. To predict how this affects prices, I develop a model of the demand for floorspace, building on the quantitative spatial models of Ahlfeldt et al. (2015) and Tsivanidis (2019). These workhorse models rely on data on commuting times and floorspace prices (or land prices) but not on floorspace quantities, which are treated as structural residuals. Accordingly, these models are typically used to evaluate the effects of changing commuting times but not to predict the effects of adjusting the amount of floorspace in a neighborhood. In contrast, the model developed in this section directly incorporates data on quantities of floorspace, which I can vary in counterfactuals. Matching the observed data on the amount of floorspace consumed in each neighborhood requires modeling two important features of reality: the high heterogeneity in workers' earnings capacity and the fact that housing is a necessity, which leads to non-homothetic preferences. A model without these features would fail to rationalize, for instance, the fact that the typical resident of the Upper East Side spends more than five times more on housing than the typical resident of the Bronx.[5]

[5]In Appendix E.1, I calibrate a simple quantitative spatial model with homogenous agents and homothetic preferences and describe why it fails to adequately predict floorspace consumption.

Workers. I partition the city into a finite set of neighborhoods populated by L workers as well as firms that produce a single final good, the numéraire.[6] Workers consume the final good and housing, and choose where to live and where to work. They are perfectly mobile within the city.

Each worker has a type θ, drawn from a distribution $F(\theta)$. Workers of type θ have a skill level $s(\theta)$, corresponding to the number of efficiency-adjusted units of labor they can provide. A worker o living in neighborhood i and working in neighborhood j has the following Stone-Geary utility:

$$U_{ijo} = \frac{B_i z_{io}^H z_{jo}^W}{d_{ij}} C_{ijo}^{1-\beta} (H_{ijo} - \underline{H}_i)^{\beta}, \tag{6}$$

Where C_{ijo} is the worker's consumption of the final good, H_{ijo} is their consumption of residential floorspace, B_i is the amenity level in neighborhood i, and $d_{ij} \geq 1$ is an iceberg commuting cost. This commuting cost is a function of the commuting time τ_{ij} between neighborhoods i and j, with $d_{ij} = \exp(\kappa \tau_{ij})$. z_{io}^H and z_{jo}^W are idiosyncratic shocks capturing workers' heterogeneous preferences for different home and work locations, respectively. I assume that these shocks are drawn independently from Fréchet distributions with shapes ε^H and ε^W—larger values of ε correspond to more homogenous preferences between workers. I assume the following timing of shocks: workers draw z_{io}^H, decide where to live, and then draw z_{io}^W and choose where to work. Finally, housing is a necessity good, and workers must consume at least $\underline{H}_i$ sq. ft of housing to live in neighborhood i. Poorer households spend a larger share of their income on housing, and that share approaches β as income grows.

Migration and congestion. As time passes, redevelopment increases the total amount of floorspace in the city, which tends to lower residential prices and increase wages. This raises residents' expected utility and causes the city's population to increase through migration. I assume that the number of workers in the city grows with their expected utility, with a migration elasticity ε_M of 2 (this corresponds to a typical estimate in literature—see, e.g., Bryan and Morten, 2019; Head and Mayer, 2021; Bordeu, 2023).

Population growth increases congestion on transportation networks. To reflect this in the model, I assume that commuting times $\boldsymbol{\tau}$ increase with the total population of the city L, with a congestion elasticity ε_C of 0.15, on the higher end of values identified in the literature (Akbar and Duranton, 2017; Couture et al., 2018; Akbar et al., 2023).

Firms. Using as inputs labor L_{Fj} and floorspace H_{Fj}, firms in each neighborhood j competitively produce a single final good that is freely traded within the city. Firms in j have access to

[6] When estimating the model, I use Neighborhood Tabulation Areas (NTAs) as a partition of NYC into 188 neighborhoods. NTAs are defined by the NYC Department of City Planning as aggregates of census tracts that approximately correspond to historical neighborhoods of the city (e.g., Chinatown, Hamilton Heights). They are mapped in Figure 4. Workers can also live and work in an additional location corresponding to the rest of the metropolitan area (see Appendix E.3 for details).

the following Cobb-Douglas production technology:

$$Y_j = A_j \left(H_{Fj}\right)^{\alpha_j} \left(L_{Fj}\right)^{1-\alpha_j}, \tag{7}$$

Where A_j is a Hicks-neutral productivity shifter and α_j governs the importance of floorspace relative to labor in production in neighborhood j. L_{Fj} is total amount of efficiency units of labor provided by workers in neighborhood j. The wage of an efficiency unit of labor provided in j is denoted by w_j.

Floorspace. The amount of residential floorspace H_{Ri} and commercial floorspace H_{Fi} available in a neighborhood is given by the sum of the (quality-adjusted) floorspace supplied by all parcels in that neighborhood. Residential floorspace can be rented at a price r_{Ri} and commercial floorspace can be rented at a price r_{Fi}. As richer households tend to own more real estate, I consider that the income coming from rents is redistributed to workers proportionally to their wage income.[7]

Endogenous productivities and amenities. New construction is likely to affect amenity and productivity levels in each neighborhood. Increasing the housing stock in a neighborhood may make it more attractive, as it will create opportunities for social interaction or lead to more public investment. However, higher density may also increase congestion in the area, decreasing its amenity value. Building commercial floorspace in a neighborhood will increase the number of people working there and make it more productive through agglomeration spillovers. Commercial activity may also directly affect amenities in a neighborhood. These spillovers (whether real or perceived) motivated the introduction of zoning in NYC, and its early planners separated commercial areas from residential ones to mitigate the externalities associated with stores and factories on nearby residents.

To account for these spillovers, I follow Ahlfeldt et al. (2015) and allow amenity and productivity levels in a neighborhood to vary with the density of residents ($\tilde{L}_{Ri}$) and jobs ($\tilde{L}_{Fi}$) there:

$$B_i = \bar{B}_i \tilde{L}_{Ri}^{\gamma_{\text{RR}}} \tilde{L}_{Fi}^{\gamma_{\text{CR}}} \qquad A_j = \bar{A}_j \tilde{L}_{Rj}^{\gamma_{\text{RC}}} \tilde{L}_{Fj}^{\gamma_{\text{CC}}}. \tag{8}$$

$\bar{B}_i$ and $\bar{A}_i$ measure the fundamental amenity and productivity attributes of neighborhoods, capturing exogenous features making an area more productive or attractive (e.g., riverside views, access to parks, or proximity to key transportation nodes). The spillover parameters γ capture agglomeration externalities—for instance, γ_{CC} captures productivity spillovers leading firms to make nearby firms more productive, and γ_{CR} measures spillovers of commercial activity on

[7] Total rental income is given by $T = \sum_i (r_{Ri} H_{Ri} + r_{Fi} H_{Fi})$, and the total income of worker o is therefore given by $I_o = w_j s(\theta_o)(1+t) \equiv w_j s(\theta_o) \left(1 + \frac{T}{L\mathbb{E}[ws(\theta)]}\right)$. The budget constraint they face is $I_o \geq C_{ijo} + r_{Ri} H_{ijo}$.

nearby residents.[8]

Spatial equilibrium. Given the set of parameters $\{\beta, \varepsilon^H, \varepsilon^W, \kappa, \gamma, \varepsilon_M, \varepsilon_C\}$ and the set of exogenous location characteristics $\{\boldsymbol{\alpha}, \bar{\boldsymbol{A}}, \bar{\boldsymbol{B}}, \underline{\boldsymbol{H}}, \boldsymbol{H_R}, \boldsymbol{H_F}\}$, I define a general equilibrium of this model as an allocation of $\{L, \boldsymbol{L_R}, \boldsymbol{L_F}, \boldsymbol{r_R}, \boldsymbol{r_F}, \boldsymbol{\tau}, \mathbf{w}, t\}$ such that workers maximize utility, firms maximize profits, amenities and productivities are determined by (8), total population is governed by the migration elasticity, commuting times are governed by the congestion elasticity, and floorspace markets clear.

5 Estimation

Parcel values. Developers make redevelopment decisions by comparing the value of existing buildings with the value of the structures that could replace them. To measure the value of buildings as a function of their characteristics, I rely on the sales data and estimate (separately for residential and commercial structures) the following hedonic model, derived from equation (3):

$$\log(p_s) = \rho^{\theta}_{1,n} + \rho^{\theta}_{2,dt} + \boldsymbol{\beta}^{\theta} \cdot \mathbf{x}_s + \nu_s, \tag{9}$$

Where p_s is the price per sq. ft observed in sale s, n is the neighborhood where the sale took place, d is an aggregate of neighborhoods, and ν_s is a price shock. For residential sales, d corresponds to a community district (which typically covers three neighborhoods), and for commercial sales, which are rarer, d is a borough. $\boldsymbol{\beta}^{\theta}$ measures the extent to which the different parcel characteristics in $\mathbf{x}_s$ are valued. $\mathbf{x}_s$ includes the building's FAR, age, landmark status, quality grade (A, B, or C), construction material, and broad type (e.g., office, retail, industrial), as well as the size of the unit(s) sold.

The value of buildings $P(h, \boldsymbol{x})$ that developers consider when making their decisions reflects the structure's use value (the discounted sum of future rents) and does not include the option value of redevelopment. Indeed, through redevelopment, developers forgo the use value of the building they are demolishing, and the new building's value only reflects its use value, as redevelopment is a terminal action. The prices that I observe in the sales data, however, likely capture both use values and option values. To ensure that my estimate of $P(h, \boldsymbol{x})$ only reflects

[8]The parametrization I retain differs from that of Ahlfeldt et al. (2015) in two ways. First, Ahlfeldt et al. (2015) only considers spillovers of firms on firms and of residents on residents, i.e., both γ_{CR} and γ_{RC} are assumed to be zero. Because spillovers across uses were a key justification for zoning, I allow these parameters to be non-zero. However, I estimate these parameters to be small, justifying Ahlfeldt et al. (2015)'s parametrization. Second, in Ahlfeldt et al. (2015), the amenity and productivity levels of a location depend not only on the density of residents and jobs in that location but also on the density of nearby locations. This is justified by the very granular scale of Ahlfeldt et al. (2015)'s analysis, where locations are city blocks. Here, the locations I consider are substantially larger, with the median neighborhood containing 135 blocks. The existing literature and my analysis below suggest that agglomeration spillovers decay quickly over space, justifying focusing on spillovers within neighborhoods in this paper.

the former, I estimate equation (9) on the sample of parcels at or above the maximum FAR limit, which have a near-zero probability of redevelopment (0.02% each year) and whose option value of redevelopment is hence negligible. I report estimated parameters for equation (9) in Appendix Table D.1. They allow me to estimate the price functions entering developers' profit function.[9]

Variable costs. The second element of the profit function is variable costs, which describe the marginal cost of increasing a parcel's FAR during redevelopment. In the absence of zoning regulations, the variable cost function could be obtained by tracing the FARs chosen by developers as a function of the floorspace prices they face.

The presence of zoning constraints complicates estimation, as the FAR chosen by a developer may no longer equalize the marginal construction cost of floorspace with its price. In many construction projects, maximum FAR limits are indeed binding, and the chosen FAR of new buildings is often close to the zoning code's maximum FAR (see Appendix Figure C.4). Despite these constraints, developers' decisions are strongly influenced by prices—Appendix Figure C.5 shows that developers facing higher floorspace prices choose higher FARs and are likelier to hit a binding FAR limit.

I use this variation to estimate the parameters of equation (4). For estimation, I make three additional assumptions. First, I parametrize α_{it}^{θ} as equal to $\alpha^0 \alpha_b^{\theta} \alpha_t$, where b is the borough in which parcel i is located. RSMeans (2020) provides a time series for construction costs as well as estimates of the relative costs of construction in different boroughs and for different building types. I use these figures to calibrate the α_b^{θ} and α_t, and estimate α^0, which captures the cost of capital for residential buildings in Manhattan in 2019. Second, I assume that the cost shock η_{it} is distributed $\mathcal{N}(0, \sigma_\eta^2)$. Finally, many buildings exceed the maximum allowed FAR because developers can exempt part of their floorspace from the limit (e.g., basements, attics, or mechanical space) or ask for a small relaxation of the FAR limit. To account for cases where developers build above the maximum FAR in the zoning code, I assume that when making redevelopment decisions, developers face the constraint $h_{\theta it} \leq \bar{h}_{\theta it} = (1+b)\bar{h}_{\theta it}^{\text{code}}$, where $\bar{h}_{\theta it}^{\text{code}}$ is the maximum allowed FAR in the zoning code, and b is a shock revealed to the developer at the same time as the cost shock and distributed according to an exponential distribution with parameter λ.

Under these assumptions, I can find for each new construction project the probability that the developer will choose any FAR as a function of the price of floorspace, the statutory zoning limit $\bar{h}_{\theta it}^{\text{code}}$, and variable cost parameters. This allows me to estimate these parameters through maximum likelihood (see Appendix D.1 for details).[10] The parameters I recover are shown in Table 1. They allow me to trace out the variable costs faced by developers as a function of the

[9] To recover expected prices in levels from the coefficients of equation (9), estimated in logs, I assume that the price shock ν is distributed $\mathcal{N}(0, \sigma_\nu^2)$, giving $\hat{p} = \exp(\hat{\rho}_{1,n}^{\theta} + \hat{\rho}_{2,dt}^{\theta} + \hat{\boldsymbol{\beta}}^{\boldsymbol{\theta}} \cdot \mathbf{x} + \hat{\sigma}_\nu^2/2)$.

[10] While the land use panel I constructed covers 2004–2022, I sometimes lack data on the characteristics of the buildings whose construction wasn't completed by 2022. Hence, I restrict my sample to 2004–2019 to estimate variable and fixed costs and simulate the city's evolution from 2019 onwards in counterfactuals.

FAR of the structure they choose to build, as well as the typical FAR developers will choose when faced with different price levels. These functions are plotted in Appendix Figure D.1.

Table 1: Variable costs parameter estimates

α^0	Baseline cost of materials	77.1	(2.2)
ζ	Capital cost share	0.516	(0.005)
λ	Exempt FAR exponential distribution parameter	1.94	(0.02)
σ_η	Cost shock variance	1.10	(0.02)

Notes: This table reports parameter estimates for the variable cost function (equation 4), indicating the cost of building structures of different FARs. It also reports the estimated value of λ, the parameter of the exponential distribution governing the extent to which developers can build over the FAR limit in the zoning code.

I estimate that in 2019, building a small structure with an FAR of 1 would cost developers \$142/sq. ft on average. A mid-rise structure, with an FAR of 4, would typically cost \$520/sq. ft, and the average construction costs for a skyscraper (with an FAR of 20) rise to \$2,357/sq. ft.

These results are consistent with available engineering estimates. RSMeans (2020) reports expected construction costs for a 1-to-3-story apartment building in Manhattan in 2019 at \$185/sq. ft. New York Building Congress (2019), which tracks the construction costs of larger structures, reports average construction costs between \$400 and \$600/sq. ft. To benchmark my cost estimates for very tall buildings, I recovered the building costs of skyscrapers built since 2000 (which had an average FAR of 23.9). I find wide variation in observed construction costs, around an average of \$1,712/sq. ft. (see Appendix Figure A.10).

The estimated capital cost share ζ, at 0.52, is similar to the estimate of 0.54 Combes et al. (2021) report for Paris, and implies average construction costs increasing almost linearly with FAR. Finally, my estimate of λ suggests that in the median construction project, developers can exempt up to 26% of the floorspace they build from counting towards the FAR limit. This order of magnitude is consistent with the types of spaces that are excluded from the FAR limit.[11]

Transaction probabilities. I observe the transaction of individual parcels in the real estate sales data, allowing me to calibrate $p_i^{\text{transaction}}$. To account for the fact that some parcels are more likely to be transacted than others (for instance, buildings with multiple ownership, like condominiums, tend to be transacted less frequently), I predict $p_i^{\text{transaction}}$ by estimating a probit regression of parcel sale dummies on parcel characteristics (see Appendix Table D.2).

Fixed costs. With estimates of the price function P and the variable cost function VC in hand, I can estimate the flow profit a developer can expect if they decide to redevelop a parcel, net of

[11] "Mechanical space, cellar space, floor space in open balconies, elevators or stair bulkheads and, in most zoning districts, floor space used for accessory parking that is located less than 23 feet above curb level" does not count towards the FAR limit (NYC Department of City Planning, 2018).

fixed costs:[12]

$$\tilde{\pi}_{it}^{\theta} = \mathbb{E}\left[\max_{h^{\text{new}} \leq \bar{h}_{\theta it}} \left\{P_{i,t}(h^{\text{new}}, \mathbf{x}_{it}^{\text{new}}) - P_{it}(h_{it}^{\text{old}}, \mathbf{x}_{it}^{\text{old}}) - VC_{it}^{\theta}(h_{it}^{\text{new}})\right\}\right].$$

Appendix Figure D.3 shows that the redevelopment probability of a parcel is essentially zero when $\tilde{\pi}$ is negative, but then increases sharply with $\tilde{\pi}$ when redevelopment becomes profitable. I can use this variation to recover estimates of fixed costs, using an extension of Rust (1987)'s nested fixed point procedure.

Indeed, given estimates of the path of profits $\tilde{\pi}$ expected by developers, parcel characteristics, sale probabilities, and values of the fixed cost parameters, I can predict the probability that the developer will take each action ($\emptyset$, R, or C) in each period. This allows me to recover fixed cost parameters through maximum likelihood estimation.

Before proceeding with estimation, I need to make assumptions about the developers' expectations about the path of variables after 2019. I assume that developers expect zoning constraints, demand parameters (productivities and amenities), and cost parameters to stay at their 2019 level. Developers have rational expectations about the path of prices and floorspace supplies after 2019: they expect floorspace quantities to evolve according to the supply model, rents to evolve according to the demand model, and prices to correspond to the sum of future rents, discounted at the rate r, which I calibrate to 5%.

Table 2 reports estimated fixed costs parameters. I find that developers wanting to build a new building on vacant land have to pay a fixed cost of $230 per sq. ft of land. Redeveloping larger structures is associated with much larger fixed costs but also profit shocks with a higher variance. Construction in denser neighborhoods is also associated with larger fixed costs, and especially more variable profit shocks. Finally, fixed costs are estimated to be 73% higher for parcels that are landmarked or in a historic district.

Model validation. The model described above aims to accurately predict developers' decisions on the extensive and intensive margins. On the extensive margin, I estimate the probability that a parcel will be redeveloped in any given year as a function of prices and zoning. On the intensive margin, I predict how much floorspace a developer will build if they proceed with redevelopment.

Appendix Figure D.4(a) evaluates model fit on the intensive margin. For redeveloped parcels, I compare the FAR chosen by developers with the model's prediction. Reassuringly, these two variables are closely correlated. Appendix Figures D.3 and D.4(b) focus on developers' choices on the extensive margin, comparing redevelopment probabilities in the model with those observed in the data. Again, the model seems to adequately predict when a parcel will be redevel-

[12] I set the characteristics $\mathbf{x}_{it}^{\text{new}}$ of new buildings to the average characteristics of buildings of type θ constructed after 2000 (except for the building's age, which is set to zero).

Table 2: Fixed costs parameter estimates

Fixed cost parameters			
δ^0	Base fixed cost	229.5	(8.8)
$\delta^{\text{demolition}}$	Demolition multiplier	2268.6	(49.0)
δ^{density}	Neighborhood density multiplier	264.0	(13.3)
$\delta^{\text{protected}}$	Protected parcels multiplier	1.73	(0.04)
Profit shock parameters			
σ_ϵ^0	Base profit shock variance	105.8	(3.3)
$\sigma_\epsilon^{\text{demolition}}$	Demolition multiplier	182.9	(6.7)
$\sigma_\epsilon^{\text{density}}$	Neighborhood density multiplier	264.8	(6.4)

Notes: This table reports parameter estimates for the fixed costs that developers must pay if they engage in redevelopment (equation 5). It also reports estimates of the σ_ϵ parameters determining the scale of the Type 1 Extreme Value distribution from which the profit shocks ϵ are drawn.

oped. The model also successfully predicts whether developers build a commercial building or a residential one, conditional on proceeding with redevelopment—see Appendix Figure D.4(c).

In counterfactuals, I will use the model to predict the evolution of NYC if zoning were relaxed in different ways. Therefore, the model will only be successful if it can correctly predict the causal effect of upzonings. To assess whether this is likely to be the case, I simulate how recently upzoned parcels would have evolved in the past decades if the upzonings had not taken place, which allows me to compute the causal effect of upzoning predicted by the model. Figure 5 compares these model-implied causal effects of upzonings to quasi-experimental estimates for the full sample of upzonings, as well as for the upzonings taking place in neighborhoods with floorspace prices above vs. below the median. The model's predictions align well with quasi-experimental results, suggesting that the model is successful at predicting the effect of zoning changes despite not being directly estimated using such rezonings.

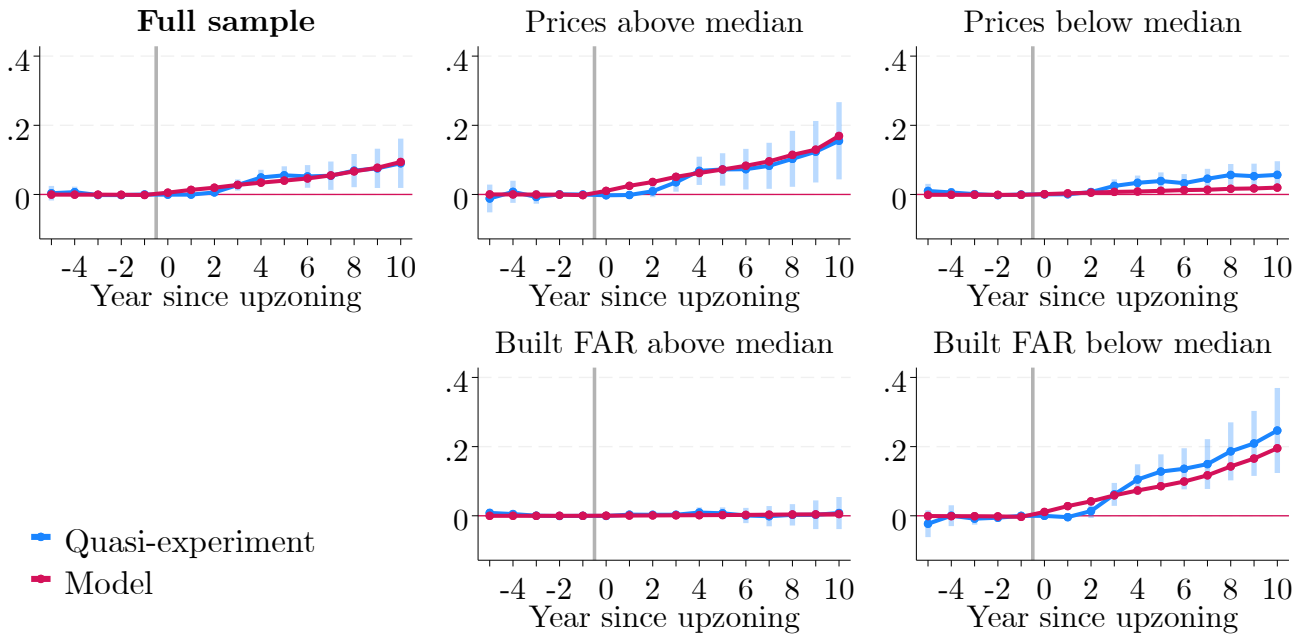
5.1 Demand

Calibration. The demand model's calibration procedure builds on the quantitative spatial economics literature (Ahlfeldt et al., 2015; Redding and Rossi-Hansberg, 2017; Tsivanidis, 2019) and is described in detail in Appendices E.2 and E.3. It can be broadly separated into four steps.

First, I calibrate the distribution of worker skills to a truncated lognormal such as to match the observed income distribution in the ACS (after excluding households with an annual income below $12,000). The ACS data also shows that the richest households in the NYC metropolitan area spend about 10% of their income on housing, and I therefore calibrate β to 0.1.

Second, I estimate the dispersion of workers' idiosyncratic preferences ε^W for work locations using data on commuting patterns in NYC. I extract from the LEHD Origin-Destination Employment Statistics (LODES) dataset the number of workers living and working in each neighborhood ($\boldsymbol{L_M}$ and $\boldsymbol{L_R}$), as well as commuting flows between them. I associate commuting flows

Figure 5: Model-implied vs. quasi-experimental effects of upzoning on built FAR



Notes: This figure compares event study coefficients of the effects of upzoning on built FAR (including floorspace under construction) with the model's predicted effects of the studied upzonings. The left panel shows results for the full sample, with the event study coefficients being the same as those reported in Figure 3. The middle (resp., right) panel shows results of this exercise for parcels in neighborhoods with floorspace prices above (resp., below) the median.

with commuting times $\boldsymbol{\tau}$, taken from Google Maps, which allows me to estimate the commuting elasticity $\kappa\varepsilon^W$ as approximately 0.044.[13] Assuming that the commuting cost parameter κ is equal to 0.01 (following Ahlfeldt et al., 2015, and Tsivanidis, 2019), I find a preference heterogeneity parameter ε^W equal to 4.4.

Third, I jointly calibrate wages $\boldsymbol{w}$, amenity levels $\boldsymbol{B}$, subsistence levels of housing $\underline{\boldsymbol{H}}$, and the dispersion of idiosyncratic preferences for home locations ε^H. Intuitively, wages can be estimated using the spatial distribution of jobs. Areas that attract many commuters are inferred to be high-wage locations, especially if they are far away from where workers live. Amenity levels can be derived from residential prices and the population distribution—neighborhoods that attract many residents despite high residential prices are inferred to provide high amenity levels. Subsistence levels of housing are estimated using data on residents' average spending on housing in each location, with higher levels of $\underline{H}$ leading to higher housing consumption. Finally, the dispersion of idiosyncratic preferences for home locations is calibrated from the observed sorting of workers across neighborhoods. In the absence of idiosyncratic preferences, workers would perfectly sort across neighborhoods as a function of their skill level, with high-skill workers systematically choosing high-amenity locations. The more dispersed idiosyncratic

[13]Following Dingel and Tintelnot (2020), I estimate the commuting elasticity using a Poisson Pseudo-Maximum Likelihood (PPML) estimator, which corrects for the bias associated with origin-destination neighborhood pairs with no commuters.

preferences are, the less pronounced such sorting patterns become. I calibrate ε^H to 2.9 to match the observed variation in average incomes across neighborhoods.

Finally, calibrated wages, along with the supply and price of commercial floorspace in each location, allow me to recover the parameters of the production function: location-specific productivities $\boldsymbol{A}$ and floorspace shares $\boldsymbol{\alpha}$.

Spillovers. The last set of parameters to calibrate are those governing agglomeration spillovers. One way to learn about these spillovers is to study the effects of new construction on the rents of nearby buildings. For example, suppose that constructing a commercial building in an area decreases the rents of nearby residential buildings. This would suggest that commercial activity induces negative externalities on nearby residents.

Therefore, to estimate γ, I isolate in the data large construction events, which I define as new buildings that increased the amount of residential or commercial floorspace within 500 ft by over 10% and by more than 20,000 sq. ft. After isolating these events, I draw 500-ft-radius disks around them and track the rent levels of preexisting buildings within these disks.[14] Data on rents is available for 290 residential construction events and 112 commercial construction events. Appendix Figure E.2 shows their location in the city, and Appendix Figure E.3 illustrates the construction of the disks used as the unit of analysis.

Comparing disks that received construction earlier vs. later allows me to evaluate the effect of new construction on rents. In Figure 6, I use the procedure of De Chaisemartin and d'Haultfoeuille (2020) to measure the effect of new construction on the quantity and price of floorspace.[15]

Upon the completion of a large structure, the amount of floorspace within a disk rises sharply. Over time, this increase in supply leads to a drop in rents in structures of the same type. However, constructing one type of floorspace has no discernable effect on the other type of floorspace.

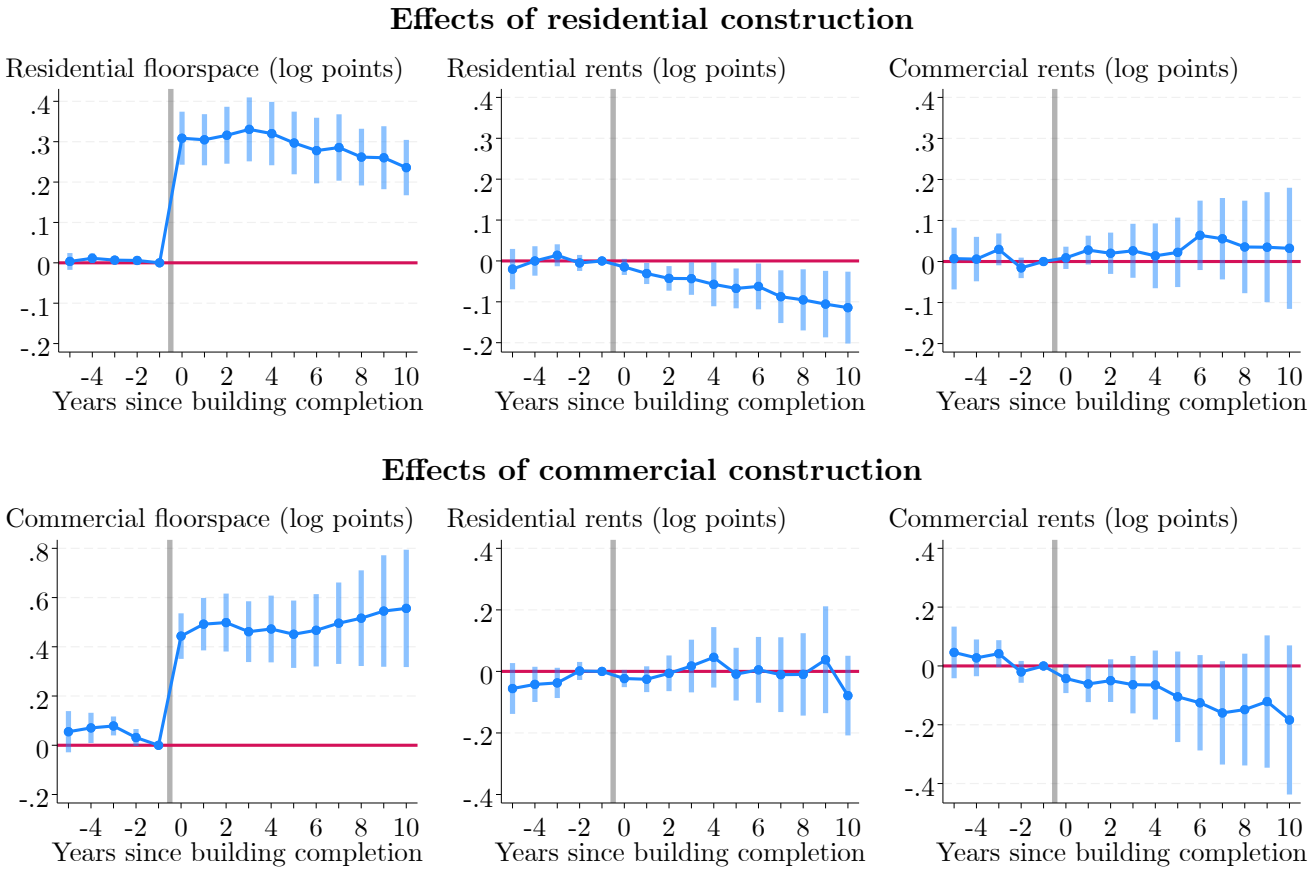
The procedure described above is similar to that proposed by Li (2022), who estimated the effects of constructing large residential towers on residential rents within 500 ft. My finding that increasing the housing stock lowers residential rents is consistent with hers, as well as those of the broader literature (Pennington, 2021; Asquith et al., 2023; Blanco, 2023). I further find that these price effects decay quickly over space—they are indistinguishable from zero beyond 500 ft (see Appendix Figure E.5), echoing previous findings (Arzaghi and Henderson, 2008; Ahlfeldt et al., 2015; Pennington, 2021; Li, 2022; Baum-Snow et al., 2024).

The event studies of Figure 6 provide us with reduced-form estimates of the elasticities of local residential/commercial rents to the local residential/commercial floorspace supply, which

[14] When computing rent levels, I exclude buildings with rent-stabilized units, as they might not appropriately reflect market conditions.

[15] In dense areas of the city, the disks used as the observation in this analysis sometimes overlap. I exclude such overlapping events from the event study, and show in Appendix Figure E.4 that the results of Figure 6 do not meaningfully change when including these overlapping disks.

Figure 6: Effects of new construction on floorspace quantities and rents



Notes: This figure shows event studies estimating the effects of new construction projects on the quantities and rent prices of floorspace in a 500-ft-radius disk around the new building. In the top panel, I study the effects of constructing a new residential building that increased the residential floorspace within 500 ft by more than 10% and by more than 20,000 sq. ft. I show effects on floorspace quantities, as well as on the residential and commercial rents of preexisting buildings. The bottom panel shows analog effects for new commercial buildings. Treatment effects are computed using the procedure of De Chaisemartin and d'Haultfoeuille (2020), comparing areas that witnessed construction earlier vs. later. In regressions, I include one observation per event and year, and cluster standard errors at the neighborhood level.

I denote by ε^{RR}, ε^{RC}, ε^{CR}, and ε^{CC}. I measure these elasticities using the average of the coefficients for the three last period of the event studies. For instance, the elasticity of residential rents to residential floorspace supply, ε^{RR}, is estimated at -0.42.

In the absence of agglomeration externalities, increasing the housing stock in an area leads to a decline in residential rents there. Negative externalities from residents on other residents ($\gamma^{RR} < 0$) will amplify the magnitude of this price response, while positive externalities ($\gamma^{RR} > 0$) will dampen it—very large positive values of γ^{RR} can even cause supply increases to lead to price increases. Increasing the residential stock in a neighborhood also affects firms there, by making it easier for them to attract workers (and inducing rises in residential rents). Positive externalities of residents on firms ($\gamma^{RC} > 0$) can lead to larger price responses.

More generally, there exists a mapping between values of γ and values of ε, which can be recovered through simulation: for any set of values for γ, I can simulate in the model the price

effects of increasing floorspace quantities. I therefore calibrate γ such as to match in the model the estimated reduced-form elasticities ε. The results of this indirect inference procedure are presented in Table 3.

Table 3: Spillover parameter estimates

Parameter	Interpretation	Calibrated value	Targeted elasticity
γ^{RR}	Effect of residents on amenities	0.12	ε^{RR} = -0.42
γ^{RC}	Effect of residents on productivity	0.03	ε^{RC} = 0.14
γ^{CR}	Effect of jobs on amenities	-0.03	ε^{CR} = -0.03
γ^{CC}	Effect of jobs on productivity	0.07	ε^{CC} = -0.28

Notes: This table reports parameter estimates for the spillover parameters γ. They are calibrated through indirect inference to match the reduced-form elasticities ε estimated using the event studies of Figure 6. ε^{XY} measures the effect of increasing the supply of type-X floorspace (residential or commercial) on type-Y rents (also residential or commercial). While the γ parameters are jointly calibrated to match all reduced-form elasticities, the value of γ^{XY} is mostly determined by the value of the corresponding reduced-form elasticity ε^{XY}, as shown in Appendix Figure E.6.

I find that amenity levels in a neighborhood increase with the number of its residents, and its productivity grows with the number of jobs there. The agglomeration elasticities I find (0.12 and 0.07) align closely with existing estimates, such as those reported by Ahlfeldt et al. (2015), which are 0.16 and 0.07, respectively. My results suggest smaller externalities across land uses.[16]

Model validation. The calibrated demand model successfully matches a set of untargeted moments of the data, as shown in Appendix Figures E.8 and E.9. First, the model accurately replicates the Engel curve for housing. Second, it successfully predicts the sorting of workers across neighborhoods, with a correlation of 0.93 between the average neighborhood income in the model and that in the ACS. Third, the model reliably predicts commuting flows between neighborhood pairs. Fourth, the calibrated shares of floorspace in production α average 0.18 (with an interquartile range of 0.14 to 0.21), matching typical estimates of the literature (Valentinyi and Herrendorf, 2008; Ahlfeldt et al., 2015). Furthermore, neighborhoods whose commercial floorspace is primarily used for manufacturing or storage uses have higher calibrated values of α than neighborhoods where it is mostly office or retail space. Fifth, the calibrated subsistence housing levels $\bar{h}$ average 224 sq. ft (with an interquartile range of 170 to 265 sq. ft), about half the size of a studio, and corresponding well to what would intuitively be the minimum quantity of housing a person could consume. Moreover, the estimated $\bar{h}$ are positively correlated with the average unit size in each neighborhood, reflecting the fact that workers tend to consume more housing in neighborhoods with larger housing units. Sixth, the model-derived amenities B are

[16]While modern-day commercial activity (in retail and office buildings) seems to have a limited impact on local amenities, this probably was not the case when zoning was introduced. Up to the 1960s, manufacturing accounted for approximately 30% of employment in NYC, and commercial activity was often associated with strong negative externalities. While new industrial construction is now rare in NYC, I find some evidence suggesting that it decreases nearby residential rents (see Appendix Figure E.7).

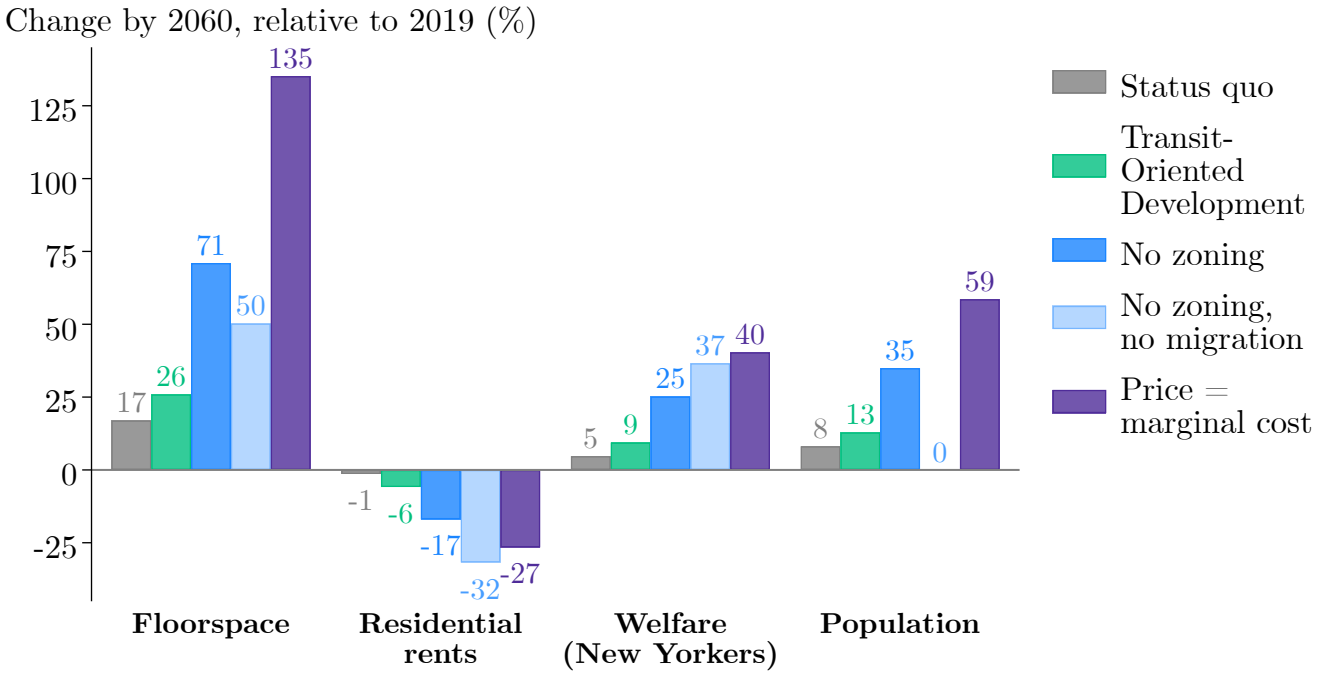
tightly correlated with residents' reported satisfaction with their neighborhood's cleanliness, access to parks, and public safety, among other amenities (see Appendix Table E.1).

6 Results

6.1 Aggregate effects of relaxing zoning regulations

With the estimated model in hand, I predict NYC's evolution until 2060 under several zoning scenarios, allowing me to evaluate the extent to which zoning influences redevelopment. As a baseline, I simulate NYC's trajectory if its zoning regulations stayed at their 2019 levels. In this status quo scenario, the city continues to grow slowly, with its floorspace supply expanding by 17% and its population increasing by 8% (see Figure 7).

Figure 7: Effects of relaxing zoning



Notes: This figure shows the predicted evolution of NYC by 2060. In the status quo scenario, zoning regulations remain at their 2019 level. In the Transit-Oriented-Development counterfactual, I upzone parcels located within 0.5 miles of a transit station (see Appendix Figure C.6). In the no zoning counterfactual, I remove all zoning restrictions (excluding parcels that are landmarked, in a historic district, or in a flood zone). In the no zoning and no migration counterfacual, I remove all zoning restrictions and set the migration elasticity to zero, such that the population of the NYC metro area stays at its 2019 level. Finally, I compare the results of my model to those of a static benchmark, in which the amount of floorspace in each neighborhood is determined by equating the price of floorspace with its marginal cost of construction, and where the price of residential floorspace is equal to that of commercial floorspace.

Transit-Oriented Development. I then evaluate the effects of an ambitious but realistic upzoning policy that increases FAR allowances around transit stations. Such policies, known as

Transit-Oriented Development, are increasingly popular around the world and are being discussed as an avenue to facilitate development in NYC under Mayor Adam's "City of Yes" plan. The Transit-Oriented Development counterfactual I consider upzones parcels within 0.25 miles of a transit station up to 4 FAR points (both for commercial and residential uses) and further increases the allowed residential FAR of parcels within 0.5 miles of a station to 3. Parcels that have higher FAR allowances at baseline are not affected by the policy change. Furthermore, I exclude landmarks, historic districts, and flood zones from the rezoning. The upzoned parcels are mostly in Brooklyn and Queens and account for 27% of NYC's developable land (see Appendix Figure C.6). I estimate that a Transit-Oriented Development upzoning of NYC would substantially increase NYC's growth, increasing the city's floorspace supply by 26% (relative to its 2019 level) and its population by 13%.

Effects of removing zoning. To assess how effectively this policy change could reap the potential benefits of relaxing zoning, I compare it to an extreme counterfactual scenario in which zoning regulations are entirely removed across the city (with the exception of parcels with landmarks and those in historic districts or flood zones). In this no-zoning counterfactual, NYC's floorspace supply grows by 71% and its population by 35%. Thus, the Transit-Oriented Development policy, which increases floorspace by 9 pp relative to the status quo, generates a boost in floorspace supply amounting to a sixth of the 54 percentage point increase of the no-zoning counterfactual.

While removing zoning substantially increases the amount of floorspace available in NYC, it only moderately reduces residential rents (by 17%). This is because the increase in floorspace supply is met with an increase in demand through migration. If there were no migration to the NYC metropolitan area as it increases its floorspace supply (for instance, because the expected utility in other parts of the country increased at the same rate as in NYC), removing zoning would lead to much greater reductions in rents and increases in welfare. This suggests that the welfare gains from relaxing zoning in NYC largely accrue to workers outside of the NYC metropolitan area.

Static benchmark. In the absence of zoning, parcels that are redeveloped reach a FAR that equalizes the construction cost of the marginal sq. ft of floorspace with its price. Hence, a natural approach to predict what NYC would look like in the long run if zoning were removed would be to compute the distribution of floorspace in the city that equalizes prices and marginal costs. I find that such a static exercise would vastly overstate the expected effects of removing zoning, predicting that NYC would more than double in size if zoning were abolished. Even if zoning exerts a strong constraint on development and drives an important gap between floorspace prices and construction costs, removing zoning could not fully close that wedge. Technological constraints, i.e., the important costs associated with increasing the amount of floorspace on a

parcel, are indeed responsible for a large fraction of the disconnect between prices and marginal costs.[17]

By assuming away technological frictions to redevelopment, a simple static model fails to capture the strong historical persistence in cities' structure. Appendix Figure F.2 illustrates the future distribution of floorspace in NYC, computed using the approach of this paper as well as a simple static benchmark. The dynamic model projects a future NYC that retains a much closer resemblance to its current structure than the static model anticipates, even in the absence of zoning. One neighborhood in which the static and dynamic approaches diverge the most is Midtown Manhattan. The dynamic model predicts that this very dense neighborhood would see its floorspace grow by 7% by 2060 if zoning were removed. In the static benchmark, where the height of buildings can frictionlessly adjust up or down, Midtown shrinks by 59% (Appendix Figure F.3).

Effects on additional outcomes. Appendix Figure F.4 illustrates additional effects of relaxing zoning regulations. Removing zoning increases the amount of commercial floorspace firms can use, leading to slightly higher wages. This new construction of commercial floorspace mostly takes place in peripheral residential neighborhoods and leads to a lower separation of land uses in the city. This tends to reduce commuting times, although this effect is offset by increased congestion caused by the city's growth. Overall, zoning's effect on commuting times appears negligible. New construction also shrinks the gap between the price of residential and commercial floorspace described in Figure 4. These price differentials, however, remain sizable even in the absence of zoning.

Use vs. bulk regulations. Zoning restricts both the use and density of new buildings. In Appendix Figure F.5, I explore the effect of both types of regulation on the city's trajectory. I find that only relaxing use limits would not meaningfully boost floorspace supply in NYC. Because most buildings in the city are near or over the FAR limit and redevelopment is usually only profitable when upward growth is allowed, new construction in the city is mainly constrained by density limits. Use limits are not irrelevant; rather, they interact with density limits, and relaxing both leads to significantly more construction than relaxing FAR limits alone.

Transition path. Appendix Figure F.6 shows the evolution of key variables under different zoning scenarios. Redevelopment slowly increases the supply of floorspace over time, leading to declines in floorspace rents and progressively decreasing the profitability of redevelopment,

[17]The large adjustment frictions associated with demolition that I find are consistent with studies of historical fires that destroyed large portions of cities (e.g., Hornbeck and Keniston, 2017; Siodla, 2015). Such studies find that burned areas become substantially denser than comparable unburned areas in the decades following the disaster. By "paying" the fixed cost associated with demolition, such fires pave the way for redevelopment, allowing for the construction of new, taller structures.

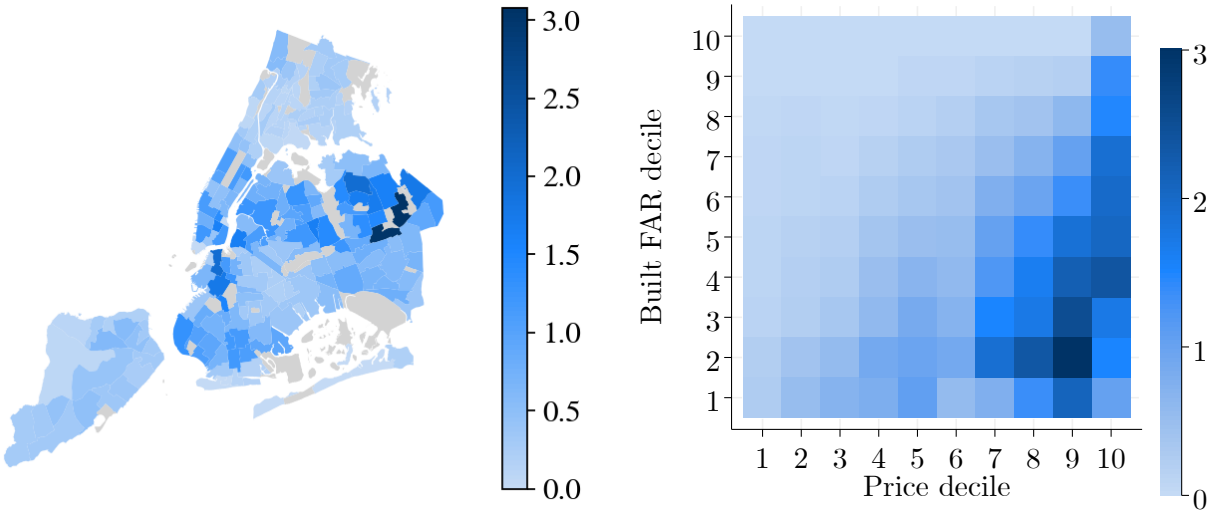
which becomes less common over time. Although relaxing zoning does not immediately reduce rents, it causes a discontinuous drop in floorspace prices because agents correctly anticipate that the policy change will lower rents in the future.

6.2 Heterogeneous effects of upzoning

Interestingly, the effects of relaxing zoning vary widely from one neighborhood to the next. Figure 8(a) shows neighborhood-level effects of removing zoning on floorspace growth by 2060. Removing zoning has limited effects in the Bronx and Midtown Manhattan, but leads to large boosts in floorspace supply in Western Brooklyn and Northern Queens.

Figure 8: Heterogeneous effects of removing zoning on additional built FAR by 2060

(a) Additional built FAR by 2060 **(b)** Split by initial FAR and price levels



Notes: The left panel shows the effect of removing zoning on the amount of floorspace built by 2060 in each neighborhood. In the right panel, I compute for each parcel the expected effect of removing zoning on the additional FAR built by 2060, and show the average effects of the policy change for each decile of floorspace price and initial parcel FAR.

These differences are easily rationalized by the basic economics of redevelopment. In Figure 8(b), I show the average effect of removing zoning on the additional FAR built by 2060 for different deciles of floorspace price and initial FAR levels. In neighborhoods where floorspace is inexpensive, redevelopment is usually unprofitable and therefore rare. The parcels that are redeveloped there tend to be vacant or minimally developed (see Appendix Figure F.7), and new buildings tend to be smaller, often below the zoning limit. Hence, zoning only moderately distorts the growth of neighborhoods where prices are low. In NYC, this suggests that upzoning most parts of the Bronx would not lead to much additional growth.[18]

[18] Appendix Figure C.7 maps floorspace prices across NYC's neighborhoods. Floorspace prices are lowest in the Bronx and highest in the lower half of Manhattan. In Appendix Figure F.8, I show the share of parcels at or above the zoning limit for different price and density levels.

Similarly, parcels that are already densely built are unlikely to be redeveloped, as the fixed costs associated with doing so are particularly high. Hence, zoning limits do not usually distort the amount of floorspace supplied there, and accordingly, my model predicts that removing zoning would only very slightly boost the growth of Midtown Manhattan despite its high floorspace prices.

The areas that would grow the most in the absence of zoning are those with high prices but low density levels, as the fixed costs of redevelopment there are low, and expected developer profits are high. This is why I find that removing zoning would boost floorspace supply the most in Western Brooklyn and Northern Queens, which have low-to-medium density levels, relatively high prices, and where most parcels are already built at or above the maximum allowed FAR.

If rezoning neighborhoods is politically costly, policymakers wishing to increase the floorspace supply in a city may want to focus their efforts on such neighborhoods where upzoning is likely to lead to important shifts in supply. This appears to have been the behavior of planners in NYC over the past decades: I find that moderately upzoning the median neighborhood by 1 FAR point would increase built FAR by about 0.03 points over the subsequent 10 years, which is substantially less than the average effect of 0.09 FAR points reported in Figure 3 (see Appendix Figure F.9). This suggests that the recently upzoned parcels were ripe for redevelopment and strongly responsive to policy changes.

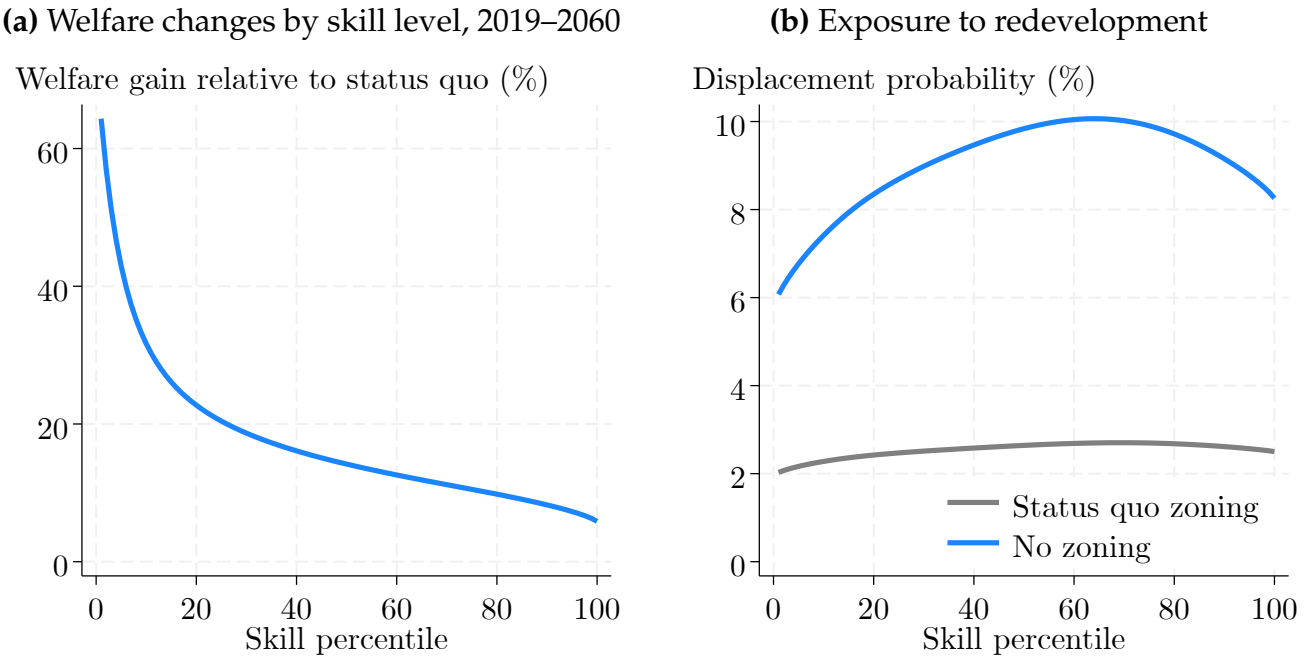
These results have two important implications for the external validity of studies of zoning changes. First, measuring the effects of past rezonings may be an imperfect guide to predicting the effects of potential future rezonings. Upzoned neighborhoods are selected based on their suitability for redevelopment, and upzoning other neighborhoods in the same city might lead to less new construction. Second, while I find that relaxing zoning constraints would substantially boost floorspace in NYC, it would likely have more limited effects in less expensive cities. In Appendix Figure F.10, I show how NYC would evolve in the coming decades if its floorspace price levels were comparable to those of Miami (with 50% lower prices) or Chicago (67% lower prices). In these scenarios, NYC's future growth would be very limited, and removing zoning would only lead to small increases in the city's overall floorspace supply.

6.3 Distributional consequences of relaxing zoning

In Panel (a) of Figure 9, I how the welfare gains of removing zoning would be distributed across the skill distribution. I find much larger gains for lower income workers, for two reasons. First, poorer workers spend a larger fraction of their income on housing, making rent decreases more consequential for them. Second, increasing housing supply in the city leads to larger rent drops in neighborhoods with lower amenity levels, where poorer households tend to live. Indeed, increasing the housing supply in high-amenity neighborhoods allows workers living in lower-amenity neighborhoods to move there and climb the neighborhood quality ladder. This lowers

housing demand in low-amenity neighborhoods, which leads to lower rents there. This mechanism, illustrated in Appendix Figure F.11, is a neighborhood analog to the "filtering" mechanism typically described in urban economics, where the supply of new high-quality housing units triggers a chain of moves along the housing quality ladder, eventually increasing the supply of affordable units (Rosenthal, 2014; Mast, 2023; Bratu, Harjunen and Saarimaa, 2023).

Figure 9: Distributional effects of removing zoning



Notes: Panel (a) shows the welfare gains from removing zoning (by 2060) for different worker types. Panel (b) displays, for each worker type, the average share of residential floorspace in their initial neighborhood that will be redeveloped. If, within a neighborhood, the redevelopment probability of a parcel is uncorrelated with the skill level of the people living there, this share corresponds to the probability that workers' initial dwelling will be demolished.

Another informative measure of the potential distributional consequences of upzoning is workers' exposure to redevelopment. While redevelopment eventually benefits households through lower rents and a more efficient allocation of land uses, they may have to move because their initial housing unit is redeveloped or otherwise suffer in the short run from the disamenities associated with construction.

Using the model, I can predict the share of the residential floorspace in each neighborhood that will be redeveloped by 2060—these shares are mapped in Appendix Figure F.12. I can then compute, for each worker type, the average share of residential floorspace in their initial neighborhood that will be redeveloped. This share, which proxies the probability that a worker will be displaced by redevelopment, is plotted in Figure 9(b). While there is overall limited heterogeneity in workers' exposure to redevelopment across types, higher-income workers tend to be more exposed to redevelopment, consistent with the fact that redevelopment is more likely

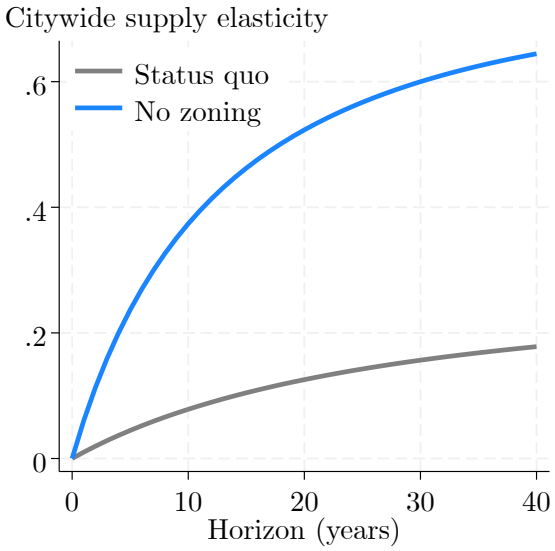
in more expensive neighborhoods.[19]

6.4 Floorspace supply elasticities

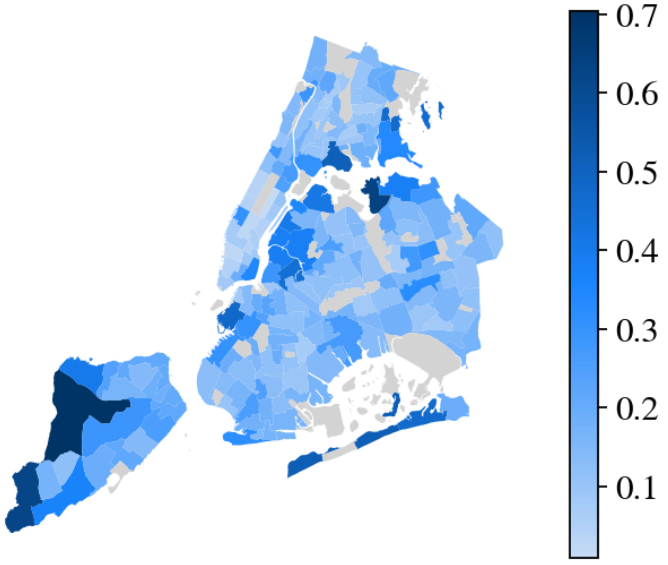
So far, I have focused on evaluating how construction would respond to changes in zoning policy, keeping fixed the aggregate demand for floorspace. The model developed in this paper allows me to conduct a symmetric exercise and estimate how construction responds to aggregate demand shocks given a set of zoning regulations, i.e., to compute floorspace supply elasticities. The model's strength is that it allows computing these elasticities on a granular scale and across various time horizons. Figure 10(a) shows NYC's overall supply elasticity at different time horizons, and Figure 10(b) shows how supply elasticities vary across neighborhoods.

Figure 10: Floorspace supply elasticities

(a) Aggregate supply elasticity at different time horizons



(b) Neighborhood-level supply elasticities (40-year time horizon, status quo zoning)



Notes: Panel (a) plots estimates of NYC's overall floorspace supply elasticity at different time horizons, both under status quo zoning and no zoning. Panel (b) maps neighborhood-level supply elasticities at a 40-year horizon, under status quo zoning.

I estimate that under current zoning regulations, NYC's overall supply elasticity reaches 0.18 at a 40-year horizon. These estimates are substantially smaller than those typically estimated at the MSA level, consistent with the results of Baum-Snow and Han (2024) who show that supply elasticities sharply decrease close to cities' central business districts. While supply elasticities are somewhat lower in denser neighborhoods, I find they are essentially determined by zoning limits. Appendix Figure F.13(a) illustrates the very strong correlation between a neighborhood's

[19]Workers at the very top of the skill distribution are less exposed to redevelopment as they tend to live in the dense neighborhoods of Manhattan, where buildings have a low probability of being redeveloped.

estimated supply elasticity and the share of the maximum FAR allowed by the zoning code (the "zoning envelope") that has already been built.

I estimate that removing zoning constraints would more than triple the city's housing supply elasticity, which would reach 0.64 over 40 years. In that counterfactual scenario, neighborhood-level supply elasticities are strongly (negatively) correlated with existing density levels—see Appendix Figure F.13(b).

Even in the absence of zoning, estimated supply elasticities are also several times lower than what a standard static quantitative spatial model would suggest. Typical parametrizations indeed imply floorspace supply elasticities multiple times larger than those I find.[20]

7 Conclusion

Cities can expand outward or upward. In the 1990s, 80% of the rapidly growing urban areas were doing so mostly via spreading. In the 2010s, this figure had shrunk to 28% (Frolking et al., 2024), as many cities were reaching a mature stage. Redevelopment is therefore becoming an increasingly important phenomenon in urban areas, allowing them to increase their floorspace supply and reallocate land to its most profitable use.

This paper provides a framework to analyze this process, leveraging the strikingly detailed and increasingly available microdata describing urban land uses. Following the city's evolution at the parcel level allows to model and separate two important barriers to redevelopment: the important costs associated with tearing down existing structures to replace them with new ones, and zoning limits on the use and size of new buildings.

I use this framework to study the effects of zoning on cities' trajectories, using NYC as a case study. I find that redevelopment is usually unprofitable in inexpensive or densely developed neighborhoods, severely hindering their ability to change regardless of the stringency of zoning. Zoning regulations primarily distort the evolution of neighborhoods with high floorspace prices and lower density. In NYC, I show that the targeted upzoning of these neighborhoods can substantially boost floorspace supply. Such zoning changes, however, largely benefit workers outside the NYC metropolitan area and low-income households, which can make reform politically challenging.

[20] Consider, for instance, a Cobb-Douglas supply function, where floorspace is produced using land and materials with a materials share of 0.75 (as, e.g., in Ahlfeldt et al., 2015). There, the floorspace supply elasticity is equal to 3 everywhere in the city.

References

Ahlfeldt, Gabriel M and Daniel P McMillen, "Tall Buildings and Land Values: Height and Construction Cost Elasticities in Chicago, 1870–2010," *Review of Economics and Statistics*, 2018, *100* (5), 861–875.

_ , Stephen J Redding, Daniel M Sturm, and Nikolaus Wolf, "The Economics of Density: Evidence From the Berlin Wall," *Econometrica*, 2015, *83* (6), 2127–2189.

Akbar, Prottoy and Gilles Duranton, "Measuring the Cost of Congestion in a Highly Congested City: Bogotá," 2017.

_ , Victor Couture, Gilles Duranton, and Adam Storeygard, "Mobility and Congestion in Urban India," *American Economic Review*, 2023, *113* (4), 1083–1111.

Allen, Treb and Costas Arkolakis, "Trade and the Topography of the Spatial Economy," *The Quarterly Journal of Economics*, 2014, *129* (3), 1085–1140.

_ , _ , and Xiangliang Li, "Optimal City Structure," *Yale University, Mimeograph*, 2015.

Almagro, Milena and Tomás Domínguez-Iino, "Location Sorting and Endogenous Amenities: Evidence from Amsterdam," Technical Report, National Bureau of Economic Research 2024.

_ , Eric Chyn, and Bryan A Stuart, "Neighborhood Revitalization and Inequality: Evidence from Chicagos Public Housing Demolitions," Technical Report, Technical Report. Working paper 2024.

Alonso, William, "Location and Land Use; Toward a General Theory of Land Rent," 1964.

Anagol, Santosh, Fernando V Ferreira, and Jonah M Rexer, "Estimating the Economic Value of Zoning Reform," Technical Report, National Bureau of Economic Research 2024.

Arzaghi, Mohammad and J Vernon Henderson, "Networking Off Madison Avenue," *The Review of Economic Studies*, 2008, *75* (4), 1011–1038.

Asquith, Brian J, Evan Mast, and Davin Reed, "Local Effects of Large New Apartment Buildings in Low-Income Areas," *Review of Economics and Statistics*, 2023, *105* (2), 359–375.

Bassett, Edward Murray, *Stores in Residence Zones* 1926.

Baum-Snow, Nathaniel and Lu Han, "The Microgeography of Housing Supply," *Journal of Political Economy*, 2024, *132* (6), 1897–1946.

_ , Nicolas Gendron-Carrier, and Ronni Pavan, "Local Productivity Spillovers," *American Economic Review*, 2024, *114* (4), 1030–1069.

Bayer, Patrick and Christopher Timmins, "On the Equilibrium Properties of Locational Sorting Models," *Journal of Urban Economics*, 2005, *57* (3), 462–477.

Blanco, Hector, "Pecuniary Effects of Public Housing Demolitions: Evidence from Chicago," *Regional Science and Urban Economics*, 2023, *98*, 103847.

Bordeu, Olivia, "Commuting Infrastructure in Fragmented Cities," 2023.

Bratu, Cristina, Oskari Harjunen, and Tuukka Saarimaa, "JUE Insight: City-Wide Effects of New Housing Supply: Evidence from Moving Chains," *Journal of Urban Economics*, 2023, *133*, 103528.

Brooks, Leah and Zachary Liscow, "Infrastructure Costs," *American Economic Journal: Applied Economics*, 2023, *15* (2), 1–30.

Bryan, Gharad and Melanie Morten, "The Aggregate Productivity Effects of Internal Migration: Evidence from Indonesia," *Journal of Political Economy*, 2019, *127* (5), 2229–2268.

Büchler, Simon and Elena Catharina Lutz, "The Local Effects of Relaxing Land Use Regulation on Housing Supply and Rents," *MIT Center for Real Estate Research Paper*, 2021, (21/18).

Burgess, Ernest W, *The Growth of the City: an Introduction to a Research Project* 1925.

Chaisemartin, Clément De and Xavier d'Haultfoeuille, "Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects," *American Economic Review*, 2020, *110* (9), 2964–2996.

Combes, Pierre-Philippe, Gilles Duranton, and Laurent Gobillon, "The Production Function for Housing: Evidence from France," *Journal of Political Economy*, 2021, *129* (10), 2766–2816.

Couture, Victor, Gilles Duranton, and Matthew A Turner, "Speed," *Review of Economics and Statistics*, 2018, *100* (4), 725–739.

Dingel, Jonathan I and Felix Tintelnot, "Spatial Economics for Granular Settings," Technical Report, National Bureau of Economic Research 2020.

Epple, Dennis, Brett Gordon, and Holger Sieg, "A New Approach to Estimating the Production Function for Housing," *American Economic Review*, 2010, *100* (3), 905–924.

Frolking, Steve, Richa Mahtta, Tom Milliman, Thomas Esch, and Karen C Seto, "Global Urban Structural Growth Shows a Profound Shift from Spreading Out to Building Up," *Nature Cities*, 2024, pp. 1–12.

Ganong, Peter and Daniel Shoag, "Why Has Regional Income Convergence in the US Declined?," *Journal of Urban Economics*, 2017, *102*, 76–90.

Gechter, Michael and Nick Tsivanidis, "Spatial Spillovers from High-Rise Developments: Evidence from the Mumbai Mills," 2023.

Glaeser, Edward and Joseph Gyourko, "The Economic Implications of Housing Supply," *Journal of Economic Perspectives*, 2018, *32* (1), 3–30.

Glaeser, Edward L, Joseph Gyourko, and Raven Saks, "Why Is Manhattan So Expensive? Regulation and the Rise in Housing Prices," *The Journal of Law and Economics*, 2005, *48* (2), 331–369.

Greaney, Brian, Andrii Parkhomenko, and Stijn Van Nieuwerburgh, "Dynamic Urban Economics," Technical Report, National Bureau of Economic Research 2025.

Gyourko, Joseph and Raven Molloy, "Regulation and Housing Supply," in "Handbook of Regional and Urban Economics," Vol. 5, Elsevier, 2015, pp. 1289–1337.

Harrison, Ballard & Allen, *Plan for Rezoning the City of New York* 1950.

Head, Keith and Thierry Mayer, "The United States of Europe: A Gravity Model Evaluation of the Four Freedoms," *Journal of Economic Perspectives*, 2021, *35* (2), 23–48.

Heblich, Stephan, Stephen J Redding, and Daniel M Sturm, "The Making of the Modern Metropolis: Evidence From London," *The Quarterly Journal of Economics*, 2020, *135* (4), 2059–2133.

Hirt, Sonia A, *Zoned in the USA*, Cornell University Press, 2015.

Hornbeck, Richard and Daniel Keniston, "Creative Destruction: Barriers to Urban Growth and the Great Boston Fire of 1872," *American Economic Review*, 2017, *107* (6), 1365–1398.

Hoyt, Homer, *The Structure and Growth of Residential Neighborhoods in American Cities*, Federal Housing Administration, 1939.

Hsiao, Allan, "Sea Level Rise and Urban Adaptation in Jakarta," 2024.

Hsieh, Chang-Tai and Enrico Moretti, "Housing Constraints and Spatial Misallocation," *American Economic Journal: Macroeconomics*, 2019, *11* (2), 1–39.

Kulka, Amrita, Aradhya Sood, and Nicholas Chiumenti, "Under the (Neighbor)Hood: Understanding Interactions Among Zoning Regulations," *Available at SSRN 4082457*, 2024.

Li, Xiaodi, "Do New Housing Units in your Backyard Raise your Rents?," *Journal of Economic Geography*, 2022, *22* (6), 1309–1352.

Liao, Hsi-Ling, "The Effect of Rezoning on Local Housing Supply and Demand: Evidence from New York City," 2022.

Lucas, Robert E and Esteban Rossi-Hansberg, "On the Internal Structure of Cities," *Econometrica*, 2002, *70* (4), 1445–1476.

Mast, Evan, "JUE Insight: The Effect of New Market-Rate Housing Construction on the Low-Income Housing Market," *Journal of Urban Economics*, 2023, *133*, 103383.

Mills, Edwin S, "An Aggregative Model of Resource Allocation in a Metropolitan Area," *The American Economic Review*, 1967, *57* (2), 197–210.

Murphy, Alvin, "A Dynamic Model of Housing Supply," *American Economic Journal: Economic Policy*, 2018, *10* (4), 243–67.

Muth, Richard F, "Cities and Housing: The Spatial Pattern of Urban Residential Land Use," 1969.

New York Building Congress, "Construction Outlook Update: New York City Construction Costs," 2019.

New York City Board of Estimate and Apportionment, "Development and Present Status of City Planning in New York City," 1914.

_ , "Final Report of the Commission on Building Districts and Restrictions to the Committee on the City Plan," 1916.

NYC Department of City Planning, *Zoning Handbook* 2018.

Ogawa, Hideaki and Masahisa Fujita, "Equilibrium Land Use Patterns in a Nonmonocentric City," *Journal of Regional Science*, 1980, *20* (4), 455–475.

Ospital, Augusto, "Urban Policy and Spatial Exposure to Environmental Risk," 2022.

Peng, Xuequan Elsie, "The Dynamics of Urban Development: Evidence from Land Use Reform in New York," 2023.

Pennington, Kate, "Building New Housing Cause Displacement? The Supply and Demand Effects of Construction in San Francisco," 2021.

Redding, Stephen J and Esteban Rossi-Hansberg, "Quantitative Spatial Economics," *Annual Review of Economics*, 2017, *9*, 21–58.

Rosenthal, Stuart S, "Are Private Markets and Filtering a Viable Source of Low-Income Housing? Estimates from a Repeat Income Model," *American Economic Review*, 2014, *104* (2), 687–706.

RSMeans, *Square Foot Costs* 2020.

Rust, John, "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher," *Econometrica*, 1987, pp. 999–1033.

Saiz, Albert, "The Geographic Determinants of Housing Supply," *The Quarterly Journal of Economics*, 2010, *125* (3), 1253–1296.

Shertzer, Allison, Tate Twinam, and Randall P Walsh, "Zoning and the Economic Geography of Cities," *Journal of Urban Economics*, 2018, *105*, 20–39.

Siodla, James, "Razing San Francisco: The 1906 Disaster as a Natural Experiment in Urban Redevelopment," *Journal of Urban Economics*, 2015, *89*, 48–61.

Soltas, Evan, "Tax Incentives and the Supply of Low-Income Housing," 2024.

Tsivanidis, Nick, "Evaluating the Impact of Urban Transit Infrastructure: Evidence From Bogota's Transmilenio," 2019.

Valentinyi, Akos and Berthold Herrendorf, "Measuring Factor Income Shares at the Sectoral Level," *Review of Economic Dynamics*, 2008, *11* (4), 820–835.

Voorhees Walker Smith & Smith, *Zoning New York City* 1958.

Appendix

A Data

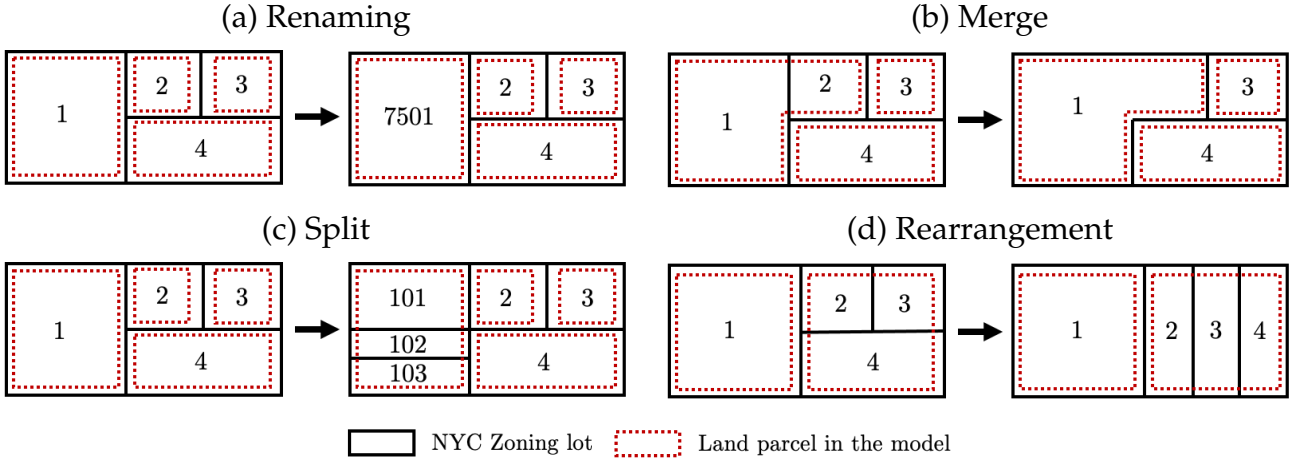
A.1 Parcel boundaries

The PLUTO dataset. Each year, the NYC Department of City Planning publishes the Primary Land Use Tax Lot Output (PLUTO) dataset, which contains detailed information about the built environment of NYC at the parcel level (called a "zoning lot"), collected for tax purposes. A zoning lot can be vacant or contain one or several buildings. In the majority of cases (about 70%), a zoning lot is associated with a single building — only 1% of the zoning lots in NYC are associated with more than 2 buildings. For each zoning lot, PLUTO contains information about building characteristics such as the year of construction, the total amount of floorspace, how it is allocated to different uses, the number of stories, etc. It also indicates the zoning district in which the lot is located, allowing me to infer the zoning regulations pertaining to it. Importantly, the PLUTO datasets are associated with GIS files indicating the boundaries of each parcel.

Defining land parcels. Zoning lots in NYC are identified by their Borough, Block, and Lot code, or BBL. It is a 10-digit number where the first digit indicates the borough in which the lot is located, the next five digits indicate a block number, and the last four digits correspond to the lot number within that block. In 2022, there were about 859,000 zoning lots in NYC, grouped into about 29,000 blocks.

While the yearly versions of the PLUTO dataset provide a high-quality repeated cross-section of the buildings in NYC, they cannot readily be used to construct a panel of land use at the lot level. This is because of numerous changes in the tax maps that took place over time. I describe in Figure A.1 the four types of changes that can be found in the data.

Figure A.1: Types of lot boundary changes



Notes: This figure shows examples of changes in lot boundaries in the PLUTO data. Solid black lines correspond to the boundaries of zoning lots, which are the most granular partition of land in NYC for which data on land use is available in any given year. Dotted red lines correspond to the boundaries of the land parcels that I use as the partition of NYC in the model. They are the smallest amount of land for which I can track land use over time.

First, there are instances in which the same piece of land will change its BBL over time (it is

"renamed"). Second, there are lot mergers — the lot stemming from the merger can take the BBL of one of the original lots involved in the merger, as is the case in panel (b), or a new BBL. Third, lots are sometimes split into several new lots. Finally, there are rarer instances in which a set of lots is merged and then split into a new set of lots. I call these "rearrangements".

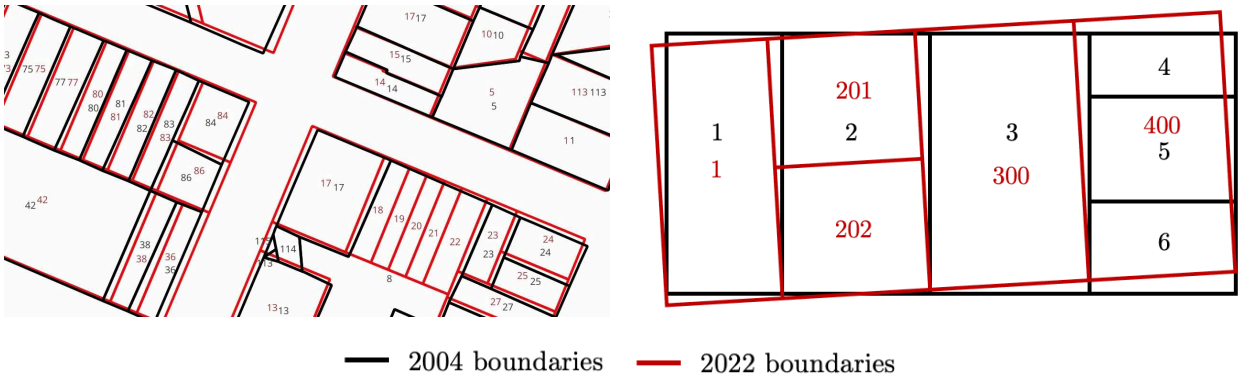
When mapping my model to the data, I consider that the land parcels that developers control are the smallest portions of land for which I can track land use over time. For instance, if zoning lots 1 and 2 eventually merge to give a new zoning lot (as in Figure A.1b), I consider the union of these zoning lots to be a land parcel in my model. I am therefore considering that zoning lots 1 and 2 are controlled by the same agent, who considers redeveloping the entire parcel made up of these two lots at each period. This assumption is justified by the fact that the merger of the zoning lots suggests that they were owned by the same individual or firm prior to the merger. In Figure A.1, I show how I define parcels for other types of lot boundary changes.

Detecting lot boundary changes. Because the same BBL can be used to identify different portions of land at different points in time, a list of lot boundary changes is necessary to build a panel of land use in NYC at the parcel level. Unfortunately, there is no dataset that documents such changes in the tax maps, and there does not seem to be any systematic rule governing changes in the numbering of lots when they occur. Sometimes, the same BBL will be used at different periods of time to identify completely different parts of a given block. To model and analyze redevelopment decisions at the parcel level, I created a database of the lot boundary changes that took place in NYC between 2004 and 2022 using tax maps for these two years. Tax maps, made available as GIS files by the city's Department of Finance, show the boundaries of all zoning lots in NYC.

The main issue that arises when comparing these GIS files is that they are not perfectly aligned. As shown in Figure A.2(a), land parcels that did not change over time will nonetheless not be exactly superposed when overlaying tax maps for different years.

Figure A.2: Detecting lot boundary changes

(a) Tax maps for 2004 and 2022 **(b)** Illustration of the algorithm



Notes: This figure illustrates the procedure that I use to detect lot boundary changes. Panel (a) overlays the tax map of a portion of the Bronx in 2004 on the 2022 tax map. Panel (b) illustrates the algorithm that I use to detect boundary changes, and is described in the text.

To detect lot boundary changes notwithstanding this issue, I used the following algorithm. After overlaying the 2004 tax map on the 2022 tax map, I considered that a 2004 zoning lot A is connected with a 2022 zoning lot B if they have a sufficiently large overlap; specifically, if $\text{area}(A \cap B) \geq 0.15 \cdot \text{area}(A)$ or $\text{area}(A \cap B) \geq 0.15 \cdot \text{area}(B)$. Then, I built a graph assembling

all of these links and extract the list of its connected sets. This procedure is illustrated in Figure A.2(b).

In that figure, lot 1 is the only one that remains unchanged over time. Among the remaining lots, overlaying the 2004 boundaries over the 2022 boundaries allows us to find the following links between lots: (2, 201), (2, 202), (3, 300), (4, 400), (5, 400), (6, 400). The connected sets in the graph defined by these edges are: {2, 201, 202}, {3, 300}, {4, 5, 6, 400}. This allows me to detect that lot 2 split into lots 201 and 202, that lot 3 was renamed to lot 300, and that lots 4, 5, and 6 merged to give lot 400. More generally, a connected set with one 2004 lot and one 2022 lot corresponds to a renaming, a connected set with several 2004 lots and one 2022 lot corresponds to a merge, a connected set with one 2004 lot and several 2022 lots corresponds to a split, and a connected set with several 2004 lots and several 2022 lots corresponds to a rearrangement. Requiring a sufficient overlap between lots to consider that they are linked limits issues caused by the imperfect alignment of tax maps.

After detecting boundary changes using this algorithm, I checked that I was able to map every piece of land included in the 2004 tax map to a piece of land in the 2022 tax map. Unfortunately, in many cases, the misalignment of tax maps was severe, which caused the algorithm to fail to detect correct boundary changes. Fortunately, however, errors in the dataset of boundary changes led the mapping from lots in 2004 to lots in 2022 to be incomplete, allowing for an easy detection of the algorithm's failures. When the algorithm could not automatically code a boundary change correctly, I compared the 2004 and 2022 tax maps and coded boundary changes manually.

I found about 21,700 lot boundary changes that took place between 2004 and 2022, affecting about 15% of the available land in NYC.[21] Most of these could be coded using the algorithm described above, but approximately 5,700 boundary changes had to be coded manually to ensure that I could track land use over time in any land parcel in NYC.

Correctly taking into account lot boundaries is essential to model redevelopment: 74% of the floorspace and 69% of the buildings constructed between 2004 and 2022 are on lots that have seen boundary changes. Hence, without a dataset documenting boundary changes, one would not be able to track which structures were replaced by the vast majority of new buildings.

A.2 Construction timelines

I built a dataset documenting all new building construction projects in NYC that started between 1989 and 2022. For each of these new buildings, I identified the year in which the project started, the year in which it was completed, and the characteristics of the new building.

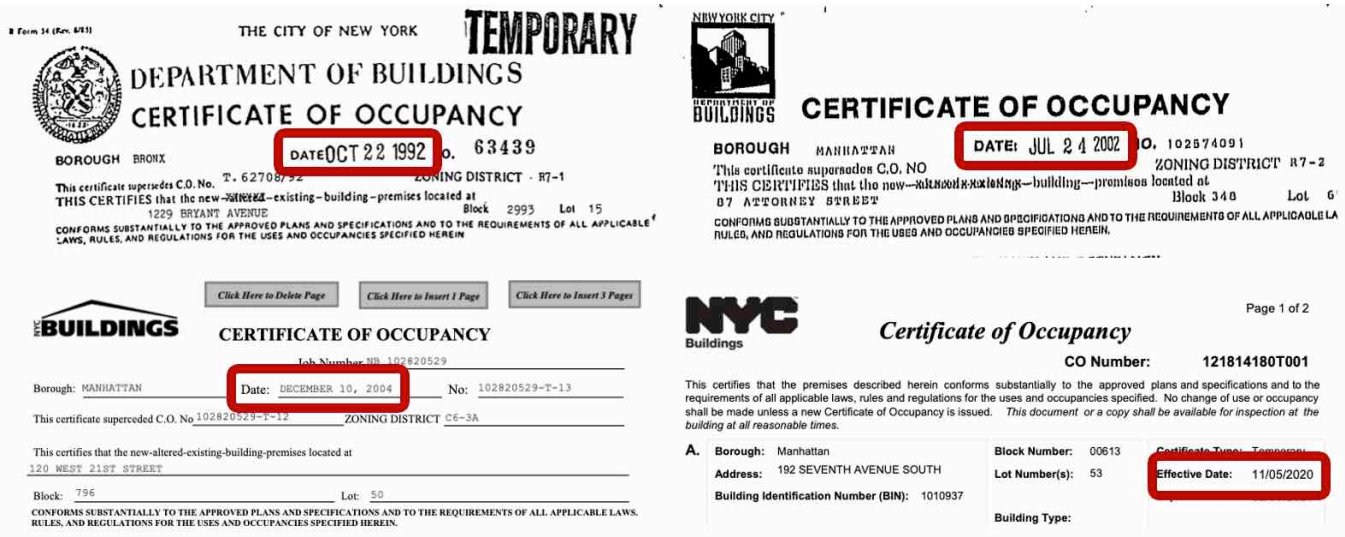
To identify the set of new building projects, I used the dataset of permits issued by the NYC Department of Buildings, which makes available the list of construction permits issued since 1989. When a developer asks for a permit to construct a new building, the project is assigned a job number. I consider the construction start year to be the year in which the first construction permit for the associated job was issued.

When construction is finalized, a Certificate of Occupancy is issued by the NYC Department of Buildings. Until this certificate is issued, nobody can legally occupy the building. Hence, I assigned as project completion year the year in which the first Certificate of Occupancy associated with the construction job is issued. I obtained data on Certificates of Occupancy from three sources. First, certificates issued after July 2012 are recorded in a dataset made available by

[21] Among the boundary changes, about 8,000 were splits, 5,800 were mergers, 4,700 were renamings, and 2,400 were rearrangements. Additionally, 430 lots appeared and 438 disappeared from the data over 2004–2022.

the NYC Department of Buildings. Second, the NYU Furman Center has compiled a dataset of Certificates of Occupancy issued for residential structures between 2000 and 2009. For buildings and time periods not covered by these two sources, I relied on the NYC Buildings Information System, which makes available scans of Certificates of Occupancy issued since the beginning of the 20th century (see Figure A.3 for examples). After scraping the relevant documents (about 250,000 in total), I used OCR software to extract the date on which they were issued.

Figure A.3: Examples of scraped Certificates of Occupancy



Notes: This figure shows the headers of four Certificates of Occupancy scraped from the NYC Buildings Information System. I extracted from these Certificates of Occupancy the date at which it was issued, highlighted in the figure.

A.3 Building characteristics

Baseline characteristics. I recovered most parcel characteristics from the PLUTO datasets, which include variables gathered by the NYC Department of Finance for property taxation purposes. These variables include a description of the parcel's use (e.g., two-family brick dwelling, stand-alone food establishment), the parcel's land area, the total amount of floorspace on the parcel and its allocation across uses (residential, office, retail, garage, storage, factory, and other), the number of buildings, (typically, one), the number of stories, the number of residential and commercial units, and year in which the parcel was developed.

Floorspace. While property tax records include data on the total amount of floorspace on a parcel, as well as for individual condominium units, they do not record the size of individual non-condominium residential units. To obtain data on the size of such units, which regularly appear in real estate sales data, I relied on StreetEasy, the NYC subsidiary of Zillow. StreetEasy's website has dedicated pages for hundreds of thousands of buildings in NYC, with information on current and past listings related to them. Importantly, these listings usually report information about units' floorspace, as shown in Appendix Figure A.5.

To extract this information, I first identified 235,399 pages on the website corresponding to NYC buildings. From each of these pages, I then extracted the building's characteristics, the list of all its units in StreetEasy's database, and past sales and rental listings associated with them. This procedure allowed me to recover information on 1.6 million residential units.

Figure A.4: Example of a redevelopment timeline



Notes: This figure illustrates the redevelopment timeline of one parcel in my dataset, located at 42 Allen Street in Manhattan's Lower East Side. At the beginning of the panel, the parcel is associated with a 2,200 sq. ft, one-story retail building that was built around 1900. This building was associated with BBL no. 1003080030. In 2013, a building permit was issued, and the parcel was redeveloped into an 11,200 sq. ft, eight-story structure. In property tax records, this new building is associated with BBL no. 1003087501.

Figure A.5: Example of scraped StreetEasy listings

OFF-MARKET DATE	UNIT	PRICE	BEDS	BATHS	SIZE
08/31/24	#PHL	**$7,735**	2 beds	2 baths	953 ft²
08/29/24	#PHJ	**$7,841**	2 beds	2 baths	984 ft²
08/19/24	#11K	**$4,252**	Studio	1 bath	502 ft²
08/02/24	#9K	**$4,252**	Studio	1 bath	502 ft²
08/02/24	#10X	**$5,400**	1 bed	1 bath	684 ft²

Notes: This figure shows an example of StreetEasy listings for the apartment building at 111 Worth Street in Manhattan. I scraped these listings to extract unit characteristics from them.

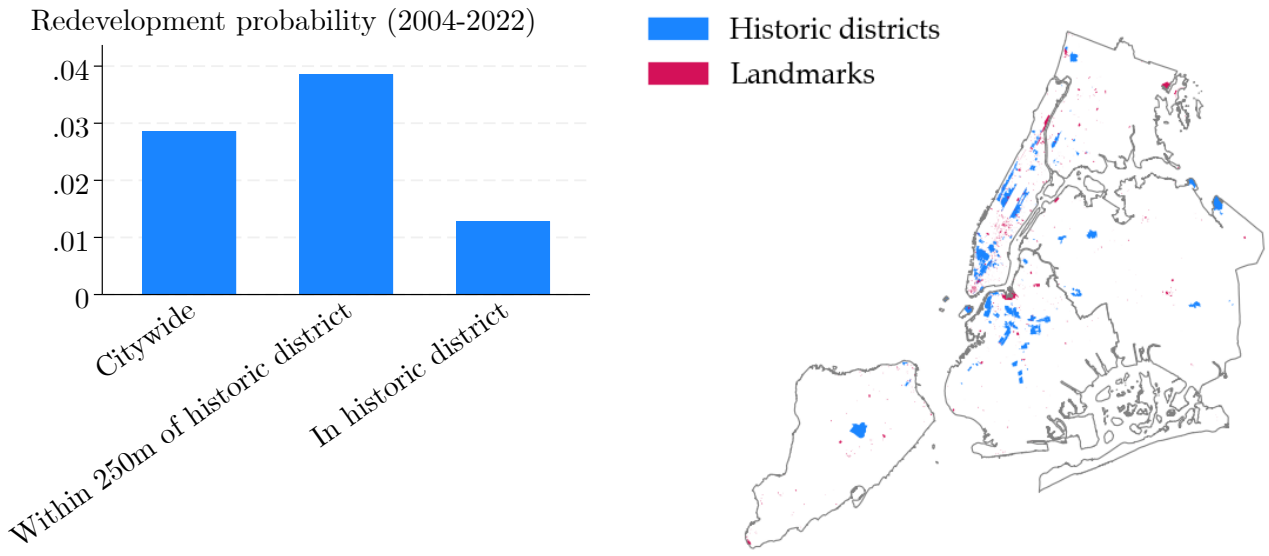
To match this dataset with NYC's administrative data, I associated buildings in StreetEasy's database with their BBL code. To do so, I first match buildings' addresses in StreetEasy with those reported in the PLUTO database. When no match was found, I relied on NYC's Department of City Planning ZoLa tool, which allows users to search for a parcel using an address.

Landmarks and historic districts. Since the establishment of NYC's Landmarks Preservation Commission in 1962, 1,462 buildings in NYC have been landmarked, and 156 areas have been designated as historic districts, covering about 32,000 parcels. To preserve the historic or architectural value of these areas, alterations of the parcels protected under one of these designations are strictly regulated and must be approved by the Landmarks Preservation Commission.

These constraints strongly reduce the probability that a parcel will be redeveloped. As shown in the left panel of Appendix Figure A.6, the redevelopment probability of a parcel in a historic district is about a third of that of a parcel in the immediate vicinity of one. The right panel

of Appendix Figure A.6 shows that these protected areas are mostly found in Manhattan and downtown Brooklyn.

Figure A.6: Parcels protected by the Landmarks Preservation Commission



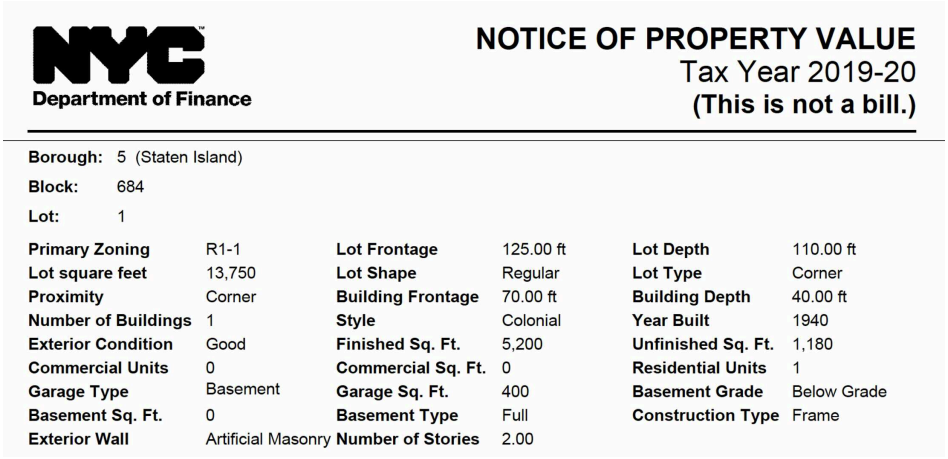
Notes: This figure illustrates the protection of some parcels by the NYC Landmarks Preservation Commission, from which I recovered the map of historic districts and landmarks. The left panel shows the probability that a parcel is redeveloped over 2004–2022 for three subsamples: the entire city, parcels that are not protected but within 250 meters of a historic district, and parcels that are in a historic district. The right panel maps the areas of NYC that are protected either as landmarks or historic districts.

Building construction materials and condition. To recover information on buildings' construction materials and condition, I relied on Notices of Property Value (NOPVs), documents that are sent yearly by the NYC Department of Finance to property owners to inform them of the latest determination of their property's taxable value. Appendix Figure A.7 shows an extract of an NOPV. These documents contain information on a range of building characteristics, including data on the building's condition and its construction materials. These variables used to determine tax assessments are not reported in the PLUTO dataset or the tax rolls published by the city. I recovered the data contained in NOPVs from three sources. First, the NYU Furman Center shared with me data from NOPVs issued between 2005 and 2015, used in Li (2022). Second, I scraped NOPVs issued between 2016 and 2023 from the NYC Department of Finance Property Tax web portal. Third, I recovered through FOIA requests to the NYC Department of Finance data used to compile the NOPVs issued between 2019 and 2023, complementing that appearing on the scraped NOPVs.

A.4 Zoning

Each year, NYC's Department of City Planning releases zoning maps in GIS format, associating parcels to their zoning districts. When rezonings occur, the Department of City Planning also releases GIS files of the boundaries of the rezoned areas. I used these files to isolate rezoned parcels.

Figure A.7: Example of a NOPV



NYC Department of Finance

NOTICE OF PROPERTY VALUE
Tax Year 2019-20
(This is not a bill.)

Borough: 5 (Staten Island)
Block: 684
Lot: 1

Primary Zoning	R1-1	**Lot Frontage**	125.00 ft	**Lot Depth**	110.00 ft
Lot square feet	13,750	**Lot Shape**	Regular	**Lot Type**	Corner
Proximity	Corner	**Building Frontage**	70.00 ft	**Building Depth**	40.00 ft
Number of Buildings	1	**Style**	Colonial	**Year Built**	1940
Exterior Condition	Good	**Finished Sq. Ft.**	5,200	**Unfinished Sq. Ft.**	1,180
Commercial Units	0	**Commercial Sq. Ft.**	0	**Residential Units**	1
Garage Type	Basement	**Garage Sq. Ft.**	400	**Basement Grade**	Below Grade
Basement Sq. Ft.	0	**Basement Type**	Full	**Construction Type**	Frame
Exterior Wall	Artificial Masonry	**Number of Stories**	2.00		

Notes: This figure shows a portion of a Notice of Property Value (NOPV) sent in 2019 to a property owner in Staten Island.

After associating parcels to zoning districts, I found in the city's zoning resolution the maximum allowed FARs for residential and commercial structures in each zoning district.

A.5 Sales data

The Property Sales Files. To compute the price of commercial and residential floorspace in NYC, I relied on the NYC Department of Finance Property Sales Files, a dataset of all property transactions that took place between 2003 and 2022. This dataset contains data on the location of properties sold and many of their characteristics (as they are relevant for tax purposes), such as the type of building, the amount of floorspace and land sold, the sale price, and the number of commercial and residential units in the property). When computing property prices, I excluded from the dataset properties that were transferred (e.g., between parents and children) and appear in the database with a sale price of $0. I further restricted the dataset to only consider sales of residential space, store space, and office space, excluding sales of industrial buildings (e.g., factories, warehouses, etc.), institutional and cultural buildings (e.g., hospitals, museums, post offices, UN buildings, airports, churches, etc.) and hotels.

Residential and commercial floorspace sales. To assess whether each sale was a sale of commercial space or residential space, I used in priority data on the number of residential and commercial units sold. If the Property Sales File records that only commercial units (resp., residential units) were sold, I consider that the transaction was a sale of commercial (resp., residential) floor space. I removed from the analysis sales that included both commercial and residential units. When data on the type of units sold (residential or commercial) is missing, I infer the type of sale with the information available on the type of building when there is no ambiguity about its use. For instance, I consider that sales of single-family homes are residential sales and that sales of retail buildings are commercial sales.

The Property Sales Files does not systematically report the total amount of floorspace and/or the amount of land sold in a transaction. To complete our dataset, I merged the Property Sales Files with the Property Assessment Database of the NYC Department of Finance, which contains data on the amount of land and floorspace associated with each property.[22] The Property

[22]The correlation between the recorded amount of floorspace (resp. land) in the Property Sales Files and that in

Assessment Database also records for a subset of properties the share of floorspace in the following categories: residential, retail, office, loft, factory, warehouse, storage, garage, and other. In the properties that I classified as commercial, over 99% of the floorspace is indeed classified as residential, and in properties classified as commercial, less than 1% of the floorspace is classified as residential (and most of the floorspace is retail or office space).[23] I excluded from the dataset properties with less than 100 sq. ft of floorspace.

I complemented the Property Sales Files with data from the PLUTO database, merging the two databases to obtain data on each property's precise location (its latitude and longitude, as well as the neighborhood in which it is located) and building characteristics.

I computed the yearly average price of commercial and residential floorspace in NYC after trimming within each year the top 5% and bottom 5% of prices (to avoid having estimates driven by outliers), and used these estimates to inflate floorspace prices using 2019 as a reference year.

To avoid having results driven by outliers, I trimmed from the data sales with a price per sq. ft above 10 times the median price per sq. ft (about $4,000) or below a tenth of the median price per sq. ft (about $40).

Overall, I have data on the sale price, amount of floorspace, and type of floorspace (residential or commercial) for about 700,000 property sales, of which a large majority (97%) are residential.

A.6 Rents

Each year, owners of income-producing properties with an assessed value exceeding $40,000 are required to file a Real Property Income and Expense (RPIE) statement and report to the NYC Department of Finance the rental income derived from their property. While the information reported in RPIE statements is confidential, the Department of Finance reports, with a two-year lag, the estimated income derived from properties in NOPVs (see Li, 2022, for additional details). This allows me to recover building-level estimates of rents.

I build on the data collected in Li (2022), corresponding to NOPVs issued between 2005 and 2015 (with, therefore, data on rents from 2003 to 2013). By scraping NOPVs issued between 2016 and 2023, I extend this dataset of rents to cover 2014–2021.[24] There are 40,259 BBLs for which rent data is available over 2004–2021, the period covered in event studies. These BBLs contain 1.4 million units out of the 4 million in the city.

A.7 Additional data sources

Smartphone location data. To explore contemporary trends in the use of commercial spaces in NYC, I gathered smartphone data from Advan, which measures the number of monthly visits to about 500,000 points of interest in the city (e.g., stores, restaurants, and banks).

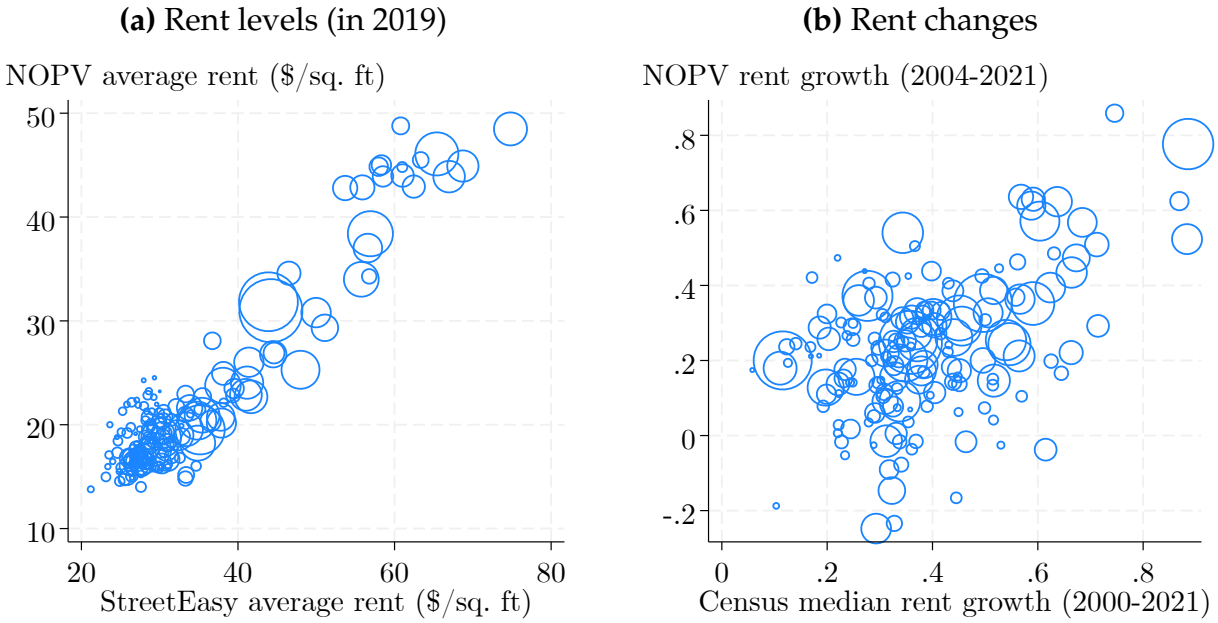
Commuting flows. I extracted data on commuting flows from the LEHD Origin-Destination Employment Statistics (LODES) database. I gathered information from the 2019 version of the dataset on each neighborhood in NYC and the broader New York–Newark–Jersey City

the Property Assessment Database is of 0.97 (resp. 0.99).

[23] Using this classification, I excluded a further subset of sales of hotels.

[24] Large fines are imposed on property owners who do not file an RPIE. However, expected rental incomes are sometimes not reported in NOPVs for some years. When rents are missing for a year within a BBL, I infer rents by linearly interpolating between the closest years for which data is available.

Figure A.8: Benchmarking of rents in NOPVs



Notes: This figure compares residential rent estimates provided in Notices of Property Values (NOPVs) with available benchmarks. In panel (a), I compare the reported rents in NOPVs for 2019 with those listed on StreetEasy for the same year. In panel (b), I compare average rent increases in NOPVs over 2004–2019 with increases in the median paid rent between the 2000 Census and the 2021 5-year ACS (changes are measured in log points and rents are adjusted for inflation). In both panels, each dot corresponds to a neighborhood, and observations are weighted by the number of BBLs with available NOPV data.

metropolitan area. Specifically, I collected data on the number of workers and jobs in each location, as well as the number of commuters traveling between any two locations.

To preserve privacy, the LEHD adds noise to the data it collects, which poses issues in some neighborhoods. For instance, about 8,000 people are reported to live in NYC's parks and cemeteries. I isolated outlier data points by computing the available residential floorspace per resident and commercial floorspace per job in each neighborhood. I adjusted the number of workers and jobs in different locations to winsorize these ratios at the 1st and 99th percentiles.

Finally, I normalized the number of workers in the NYC metropolitan area to its 2019 population of 19.2 million. Hence, expected incomes and floorspace consumptions are measured in per capita terms, facilitating their interpretation.

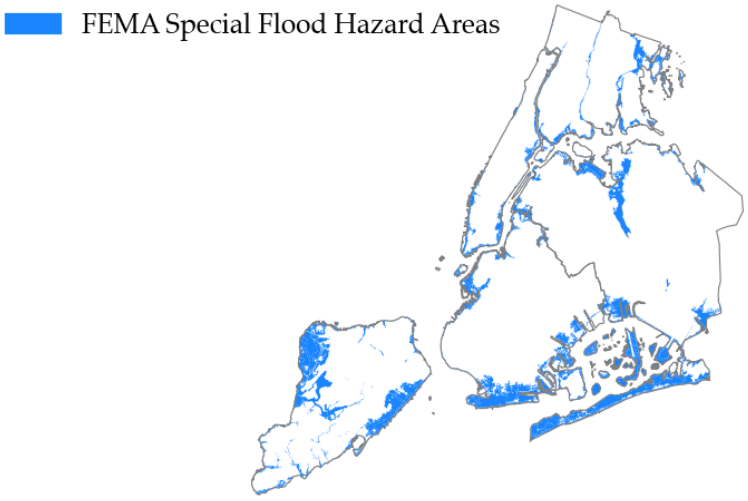
Commuting times. I gathered the travel time between any two locations using public transit or a car from the Google Maps API. Travel times are measured for a typical Monday at 8 a.m. and include traffic-induced slowdowns. I set the travel time between neighborhoods in NYC and the rest of the metropolitan area to an hour.

Train stops. I retrieved from the Baruch College Newman Library GIS Lab the locations of the train stations of the NYC Subway, the Staten Island Railway (SIR), the Long Island Rail Road (LIRR), the Metro-North Railroad, and the Port Authority Trans-Hudson (PATH).

Flood maps. I retrieved from FEMA the parts of NYC designated as Special Flood Hazard Areas, defined as areas that have a probability of being flooded in a given year exceeding 1%.

These flood zones are shown in Appendix Figure A.9.

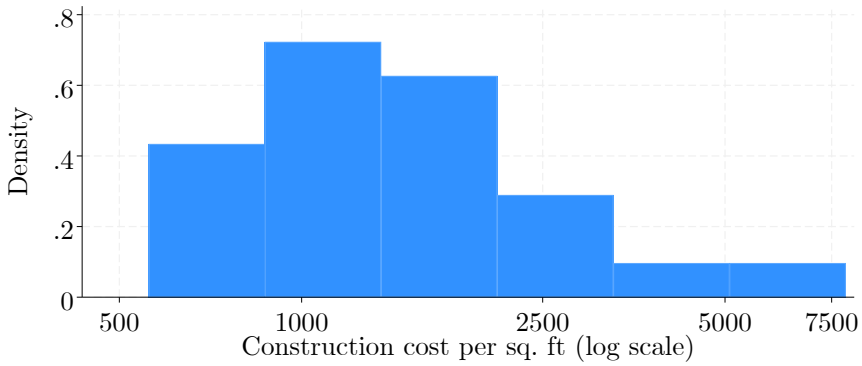
Figure A.9: Flood zones



Notes: This figure shows high-risk flood zoned in NYC, i.e., the parts of the city that are designated by FEMA as having a yearly flood probability exceeding 1%.

Skyscraper construction costs. To benchmark the construction costs of very tall structures, I gathered the list of buildings constructed after 2000 that exceed 650 ft in height and recovered their construction costs from newspapers (primarily the New York Times), specialized websites (e.g., The Real Deal, Commercial Observer, NY YIMBY), and Wikipedia pages.[25] I recovered cost estimates for 47 buildings, which have an average FAR of 23.9 and whose construction cost averaged $1,712 per sq. ft (with much variance across skyscrapers—the standard deviation of these construction costs is $1,365 per sq. ft—see Appendix Figure A.10).

Figure A.10: Skyscraper construction costs



Notes: This figure shows the distribution of observed construction costs per sq. ft for skycrapers built after 2000. Construction costs have been adjusted to be interpreted in 2019 dollar terms.

[25]To make reported construction costs comparable over time, I adjusted them using the construction price indices reported by RSMeans (2020).

B Historical and Policy Context

B.1 A short history of zoning in NYC

Early planning and the 1916 zoning resolution. In the 19th and early 20th centuries, NYC grew tremendously as it industrialized and became a major entry point for immigrants to the United States. To fight unsanitary conditions in the city's poorest neighborhoods, New York State passed legislation in 1867, 1879, and 1901 (the "Tenement Acts") to ensure that newly built dwellings provided sufficient fire protection, sanitation, ventilation, and light to their residents.

To more thoroughly regulate urban form, NYC adopted in 1916 a comprehensive zoning ordinance, becoming the first American city to do so. The motivations of the architects of the 1916 resolution are summarized in the report they submitted with the rules they proposed (New York City Board of Estimate and Apportionment, 1916). These early urban planners were, as the designers of the Tenement Acts, chiefly concerned with the effects of uncontrolled development on public health. Density limits, enacted through mandatory setbacks and height limits, would increase buildings' access to light and air, curtailing the spread of disease. A key concern was tuberculosis, which killed thousands of people each year in the city. Density limits were also justified as a means to limit the potential damage caused by fires.

On top of imposing density limits, the new zoning resolution favored the separation of residential and commercial land uses by restricting the areas where the construction of commercial buildings could take place. Such regulations aimed to segregate land uses perceived as incompatible. One reason for incompatibility is pollution and noise, which justifies the separation of factories from dwellings. Interestingly, the planners of the 1916 resolution also largely advocated for the separation of land uses on moral grounds. The presence of stores where "much loafing is done" was believed to have "a very undesirable moral effect on the children living nearby" because of the "sordid atmosphere of the ordinary business street." Stores were harshly denounced as "invading residence streets" and mixed-use buildings as "undesirable as homes and unprofitable as stores." Edward Bassett, who wrote the first comprehensive zoning ordinance of NYC, assessed that "Stores with families above should be relegated to the dark ages of the past. The play spaces of small children ought not to be near fruits and vegetables for sale. [...] One of the best tendencies of zoning is to make business streets business only and residence streets residences only" (Bassett, 1926, as cited by Hirt, 2015).

The 1961 zoning resolution. By the end of the 1940s, support for a rezoning of New York had grown enough for the city to commission the drafting of a new zoning resolution. A first plan was proposed in 1950 by the architectural consulting firm Harrison, Ballard & Allen. While this plan was rejected by the city, it served as a foundation for the work of another consulting firm, Voorhees Walker Smith & Smith, which drafted in 1958 the zoning resolution that would eventually be adopted in 1961.

One key motivation for overhauling zoning regulations was the perceived ineffectiveness of existing rules. When NYC's initial zoning ordinance was passed, the U.S. Supreme Court had not yet confirmed the constitutionality of municipal zoning, and the 1916 resolution code only modestly restricted development. Urban planners had computed that if NYC were developed to the full extent allowed by the 1916 ordinance, it would provide dwellings for 70 million residents and workspace for 300 million employees (Harrison, Ballard & Allen, 1950). The architects of the new zoning resolution, therefore, tightened restrictions to make them able to substantially affect urban form.

The economic and public health challenges that the city faced had also substantially changed since the turn of the century. Tuberculosis had been brought under control in the United States, and improved building codes had substantially reduced fire risk. Traffic congestion and the lack of parking space, however, had become a major issue as the automobile became a primary mode of transportation. In the post-war era, small neighborhood shops were closing throughout NYC as they faced increasing competition from supermarkets, leading to vacant storefronts and blighted neighborhoods.

To draft a new zoning code, the mid-20th-century urban planners predicted how the city's needs for residential and commercial space would grow until 1970. They then assigned density and use limits to different areas of NYC to guide this growth in the direction they viewed as most beneficial to the city.

The new zoning code would promote "tower in the park" developments, consisting of high-rise buildings surrounded by green spaces and parking lots. These structures, popularized by Le Corbusier, were particularly popular in the 1960s as NYC's zoning resolution was drafted. They were seen as an efficient way to provide ample light, air, open space, privacy, and parking for city residents. To favor the construction of such developments, planners adopted maximum FAR limits as the primary lever to regulate land uses in the city instead of the previous height limits and mandated setbacks, which favored buildings fronting sidewalks.

Urbanistic trends of the 1960s also favored the separation of commercial spaces from residential ones. The new zoning resolution, therefore, reduced the amount of land zoned for commercial uses in residential neighborhoods, and substantial portions of the city would also be zoned exclusively for commercial uses, particularly for manufacturing.[26]

Zoning in the 21st century. Views of what constitutes desirable land use patterns have profoundly changed since the 1960s. Tower-in-the-park developments have largely fallen out of favor in urban planning, and contemporary approaches favor mixed-use developments that dedicate less space to parking and integrate commercial and residential spaces. Recent neighborhood rezonings in NYC have favored such mixed-use spaces.

As housing costs sharply increased in NYC, city planners have also leveraged zoning as a means to increase the supply of affordable housing. In 2016, NYC adopted Mandatory Inclusionary Housing zoning requirements, which mandate new developments in recently rezoned areas to allocate a share of newly built floorspace for affordable housing.

Another important challenge facing the city is the increased risk of catastrophic weather events due to climate change. In 2012, Hurricane Sandy severely hit the city, causing an estimated 19 billion dollars of damages. In the aftermath of this disaster, NYC adapted zoning to increase the resilience of flood-prone areas.

[26]The abundant provision of land for manufacturing purposes was justified in the following way in Harrison, Ballard & Allen (1950): as NYC "does not raise its own food, nor mine [its] coal [...], the workers of New York City must fabricate goods or perform services that will pay for these indispensable imports of food and raw materials. [...] There is little opportunity to increase materially the employment in wholesale trade, in the retail stores, in finance, insurance [...] due to the fact that [the city] already has the lion's share in these activities. Factory production [will therefore] give much of the basic employment needed to sustain the City's economy." This reasoning was natural: for the preceding century, NYC's economy had largely thrived because of a dynamic manufacturing sector, which provided approximately 30% of jobs in the city. Planners assumed this would continue to be the case in the future. Yet, events took a different course: manufacturing started to decline precipitously in the 1960s, and it only provides a residual share of total employment in the city today. This decline, however, was more than compensated by the expansion of the service sector.

The restrictiveness, persistence, and gradual complexification of zoning. Zoning regulations in NYC became much more restrictive, persistent, and complex than the planners who designed them had anticipated them to be. While the 1961 zoning rules were much more stringent than the ones NYC had adopted in 1916, the planners who designed them did not view them as imposing tight constraints on the city's growth. Urbanists at the time considered that NYC would not substantially grow in the foreseeable future—Harrison, Ballard & Allen (1950) indicated, for instance, that "a population of about 8,600,000, predicted to be reached by 1970, should be regarded at the probable ultimate, i.e., stabilized, population of the city." The final zoning plan, "based on the concept of a city that can grow to a population of some 11,000,000 persons," was viewed as generously allowing new development (Voorhees Walker Smith & Smith, 1958). Even today, a naive observer of NYC's zoning ordinance might conclude that it does not restrain much NYC's growth. Indeed, two-thirds of the parcels in NYC have a built FAR below the maximum allowed by zoning, and if these parcels were built up to their FAR limit, NYC would have 69% more floorspace. This paper shows that despite this apparent large potential for growth under current regulations, the large fixed costs associated with redevelopment make zoning a strong constraint.

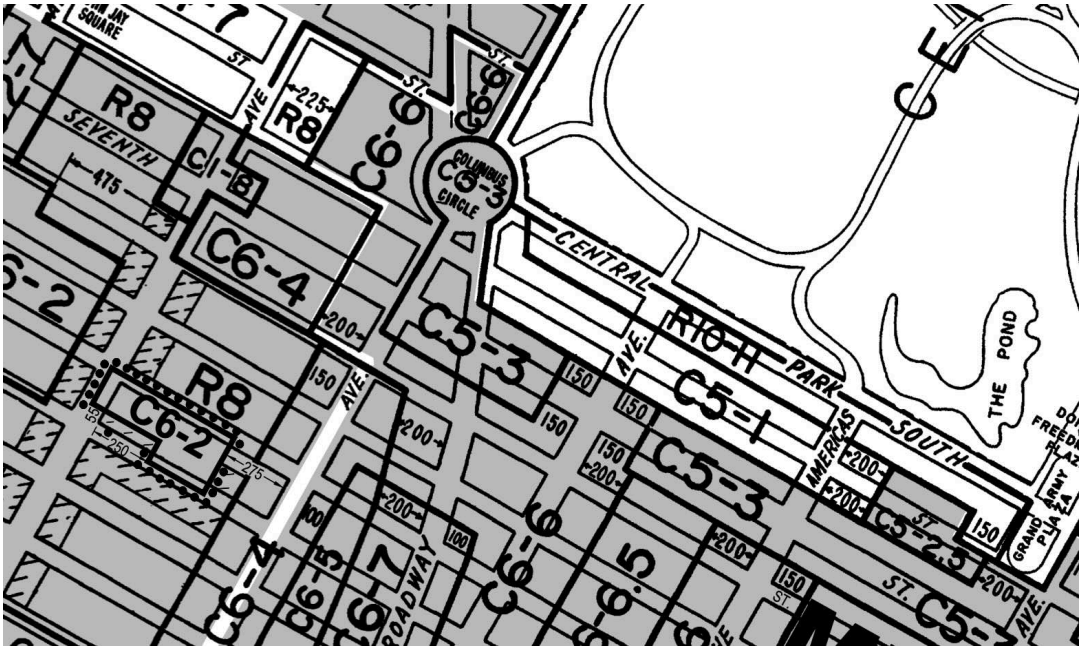
The designers of NYC's zoning resolutions never intended these planning instruments to be final or even inflexible. The crafters of the 1916 zoning ordinance insisted that "no limit can be set to the growth and expansion of the city. [...] We cannot adopt a plan and make that the Procrustean mold for all future time. The plan must develop and change with the advance of civilization." (New York City Board of Estimate and Apportionment, 1914). Similarly, the planners who first drafted the 1961 zoning resolution focused on the city's short-to-medium-run evolution and designed rules that would accompany the city's growth until 1970 (Harrison, Ballard & Allen, 1950), anticipating that zoning regulations would be adjusted in the future to guide longer-term changes in the city. While most neighborhoods of NYC have been rezoned in some way since 1961, these policy changes have typically been marginal, leading the overall zoning patterns in the city to be remarkably persistent, as shown in Appendix Figure B.4. Such persistence is not unique to NYC, as shown by Shertzer et al. (2018)'s case study of Chicago.

One of the aims of the 1961 resolution was to make zoning rules easier to understand. In part because of the 2,500 amendments that had been added to the 1916 ordinance in the decades following its approval, the constraints imposed on construction were complex and opaque. While the 1961 resolution simplified many existing rules, the progressive amendment of zoning rules did not stop, leading to gradual complexification. Over 60 years, the length of the zoning resolution grew tenfold, reaching 3,400 pages as of 2024.

The crafters of the 1961 resolution knew that the regulatory overhaul they proposed "[would] meet opposition—from those who resist change in general, from those who feel that they have a vested interest in the present resolution or who can exploit its weaknesses, and from those who honestly believe that continued piecemeal amendment is the best way to accomplish our ends" (Voorhees Walker Smith & Smith, 1958). If anything, zoning changes have become more politically contested over time, as citizens have become more involved in regulatory processes (Brooks and Liscow, 2023). Proposed upzonings—in NYC as well as in other American cities—are routinely protested against, usually because they are perceived as driving gentrification, changing neighborhoods' character, or primarily benefiting developers instead of residents. However, the recent housing affordability crisis has sparked movements like YIMBY (Yes In My Back Yard), which advocate for reforms broadly relaxing zoning regulations.

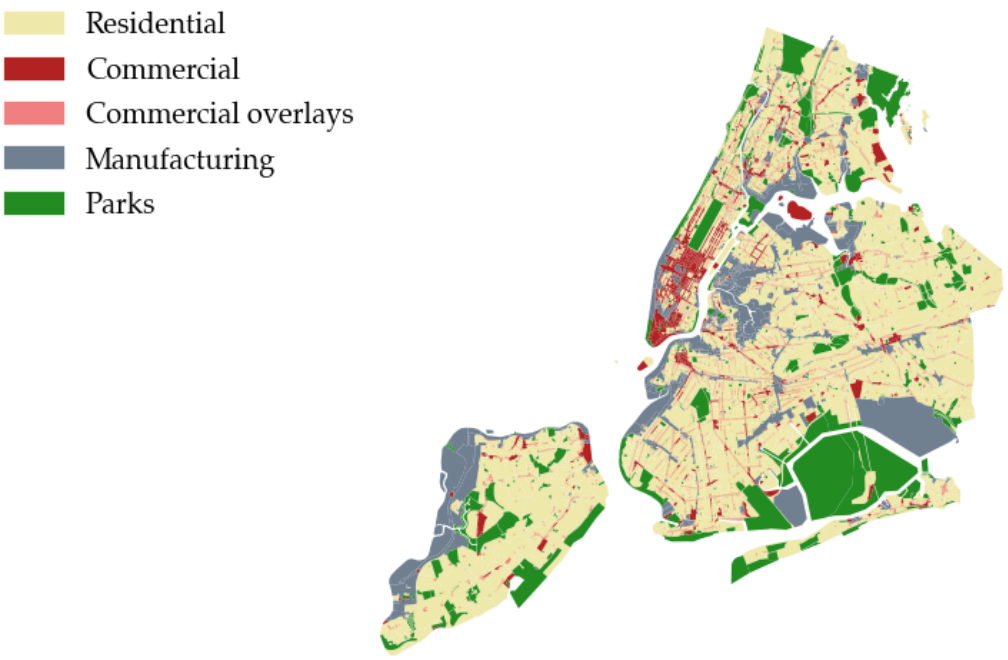
B.2 Additional figures

Figure B.1: NYC's zoning map around Columbus Circle



Notes: This figure maps zoning districts in the area surrounding Columbus Circle. Each parcel in NYC is associated with a zoning district (e.g., R8, C5-1, C6-2). Zoning districts, in turn, are associated with restrictions on the characteristics of newly built structures.

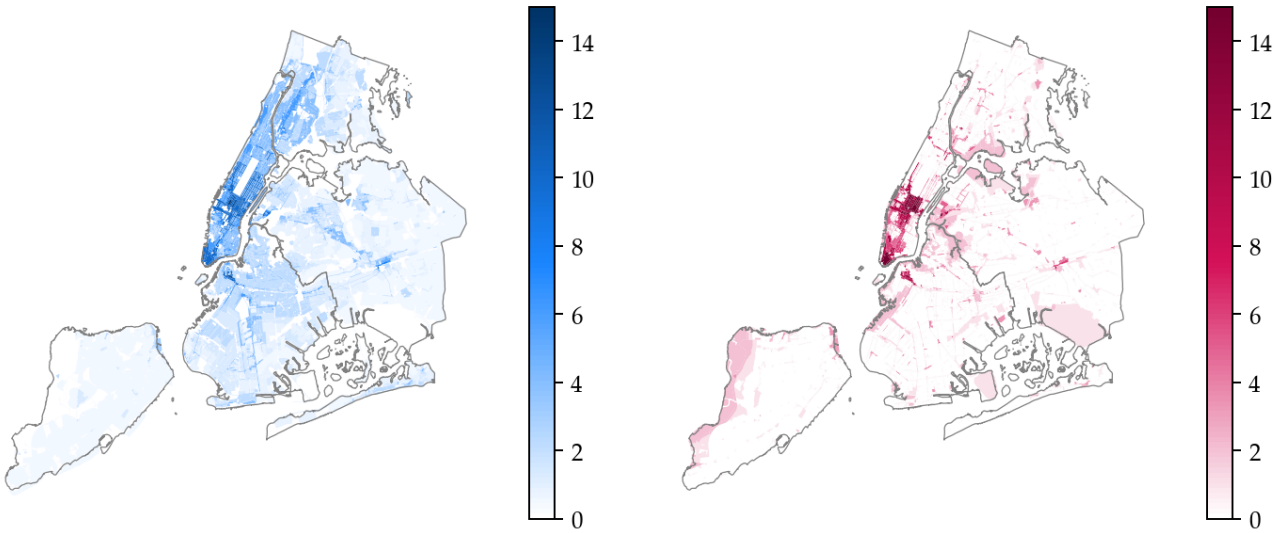
Figure B.2: Use districts in NYC's zoning map



Notes: This figure shows land use districts in NYC's zoning map. Residential districts only allow residential uses, commercial districts allow both residential and commercial uses, and manufacturing districts only allow commercial uses. Commercial overlays correspond to zones in residential areas that allow some commercial uses (e.g., stores on the first floor of buildings).

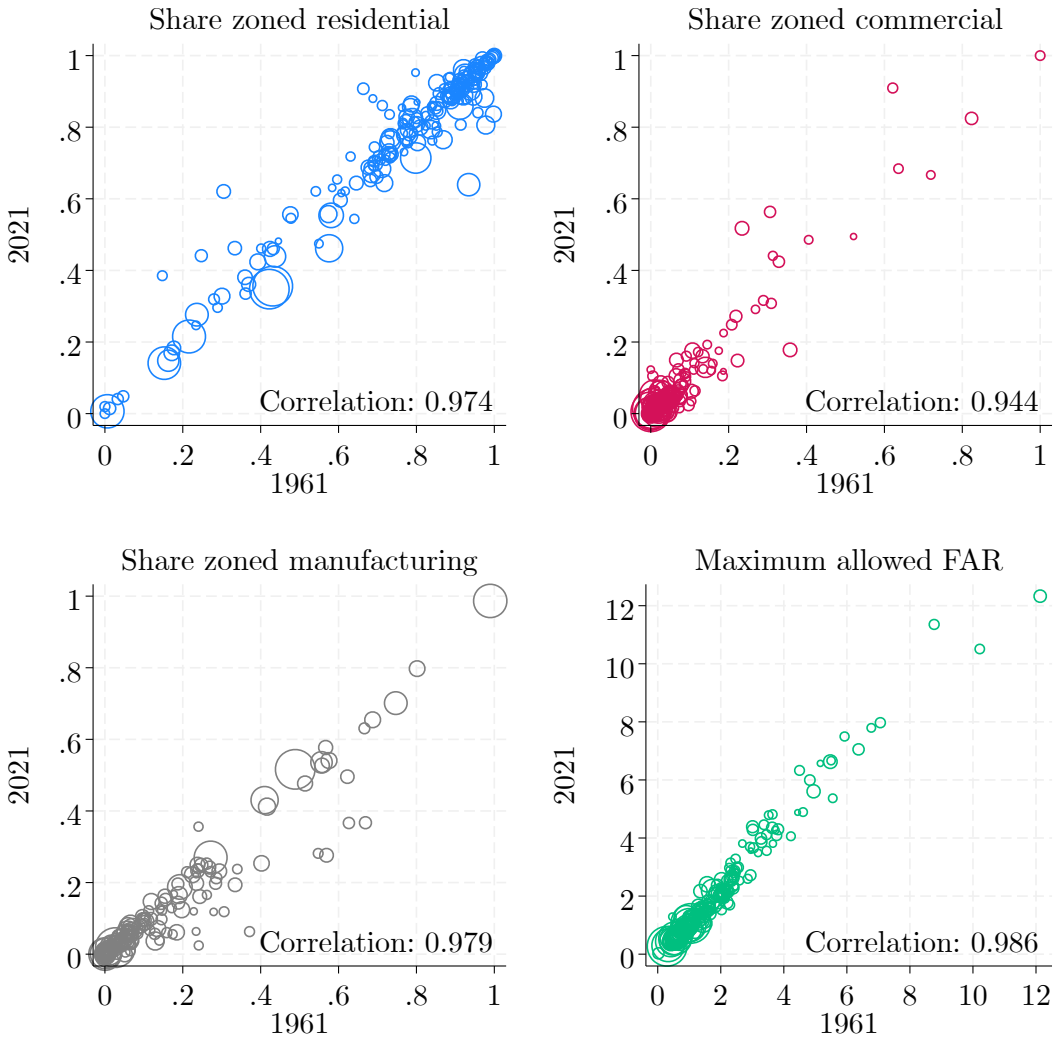
Figure B.3: FAR allowances in NYC's zoning map

(a) Residential uses **(b)** Commercial uses



Notes: This figure maps the maximum FAR allowances in NYC for residential and commercial uses.

Figure B.4: Persistence of zoning regulations

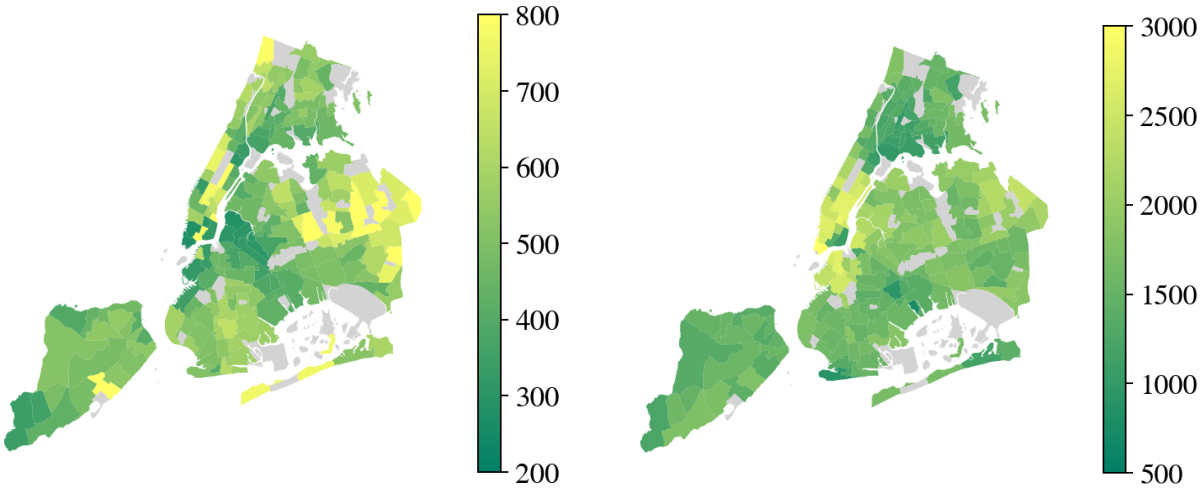


Notes: This figure shows the persistence of zoning regulations by comparing the zoning maps of 1961 and 2021. I plot for each neighborhood in both years the share of land that is in a residential, commercial, or manufacturing zone, as well as the average maximum allowed FAR. Neighborhoods are weighted by their land area.

Figure B.5: Central neighborhoods of NYC have become more attractive to residents since the 1950s

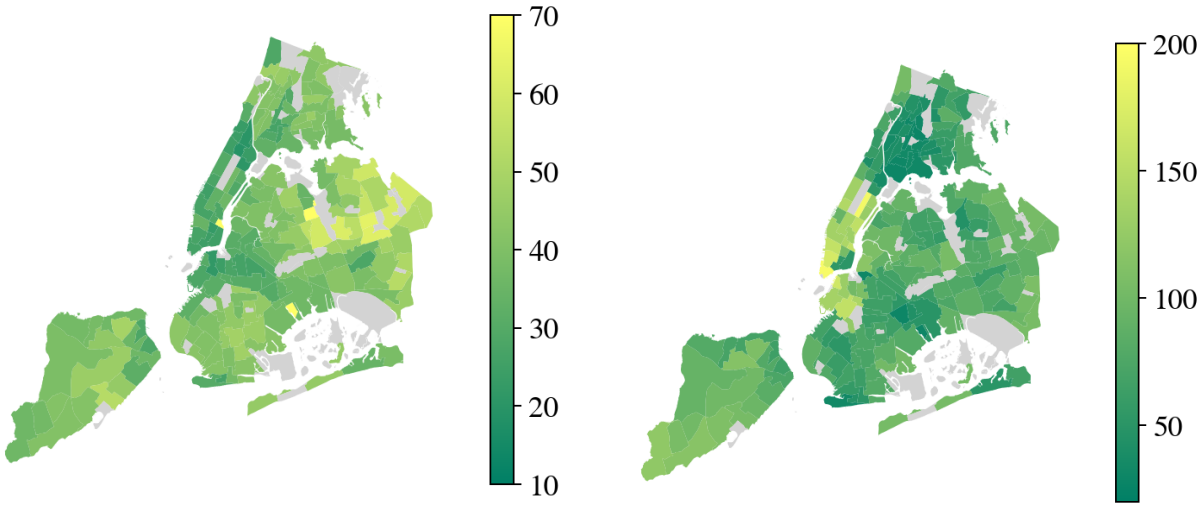
Panel A: Median residential rents (dollars per month)

(a) 1950 **(b)** 2019



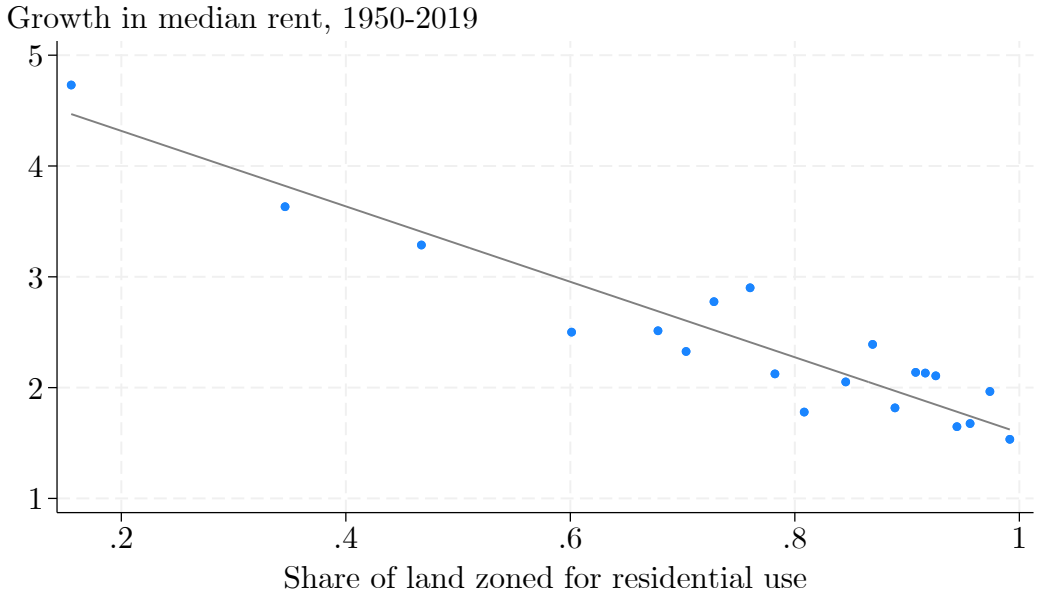
Panel B: Median income (thousands of dollars per year)

(a) 1950 **(b)** 2019



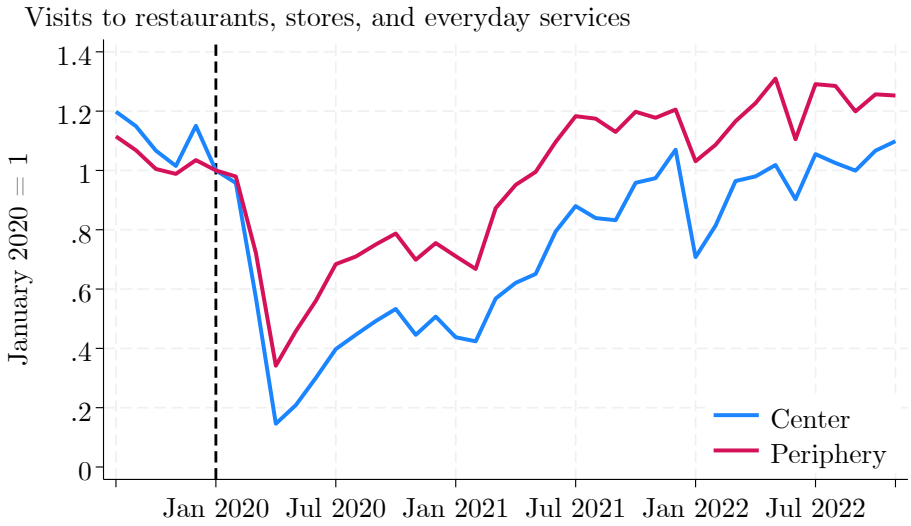
Notes: This figure compares two measures of neighborhood attractiveness in 1950 and 2019. Panel A reports median monthly rents for rented dwelling units, and Panel B reports median household incomes. Data from 1950 is taken from the decennial census, while data for 2019 is taken from the 2017-2021 five-year ACS.

Figure B.6: Since 1950, residential rents grew faster in areas of NYC zoned for commercial purposes



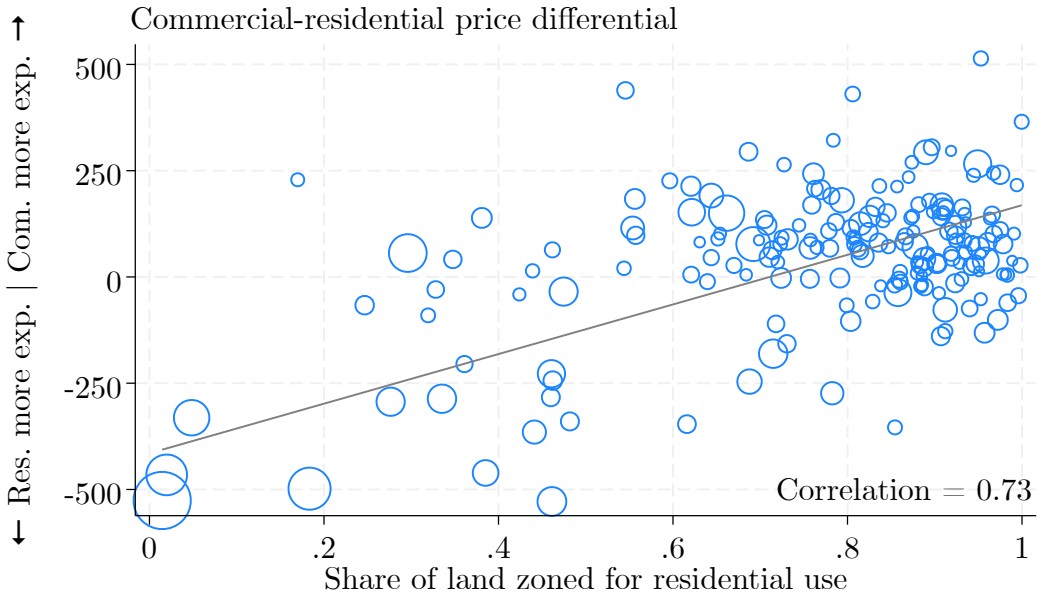
Notes: This figure illustrates the negative correlation between neighborhood-level changes in median residential rents between 1950 and 2019 (shown in Appendix Figure B.5) and the share of land that was zoned for residential purposes in 1961.

Figure B.7: Visits to commercial locations following the COVID-19 pandemic



Notes: This figure shows the evolution of the number of recorded visits to restaurants, stores, and everyday services (e.g., dry cleaners, banks) in the aftermath of the COVID-19 pandemic. I extract these trends from smartphone data provided by Advan, which measures the number of monthly visits to about 500,000 points of interest in NYC. I normalize the total number of visits in January 2020 to one and show separate trends for the center and periphery of the city. The set of central neighborhoods includes all neighborhoods in the lower half of Manhattan (below Central Park) and the neighborhood tabulation area including Downtown Brooklyn. All other neighborhoods in the city are classified as peripheral.

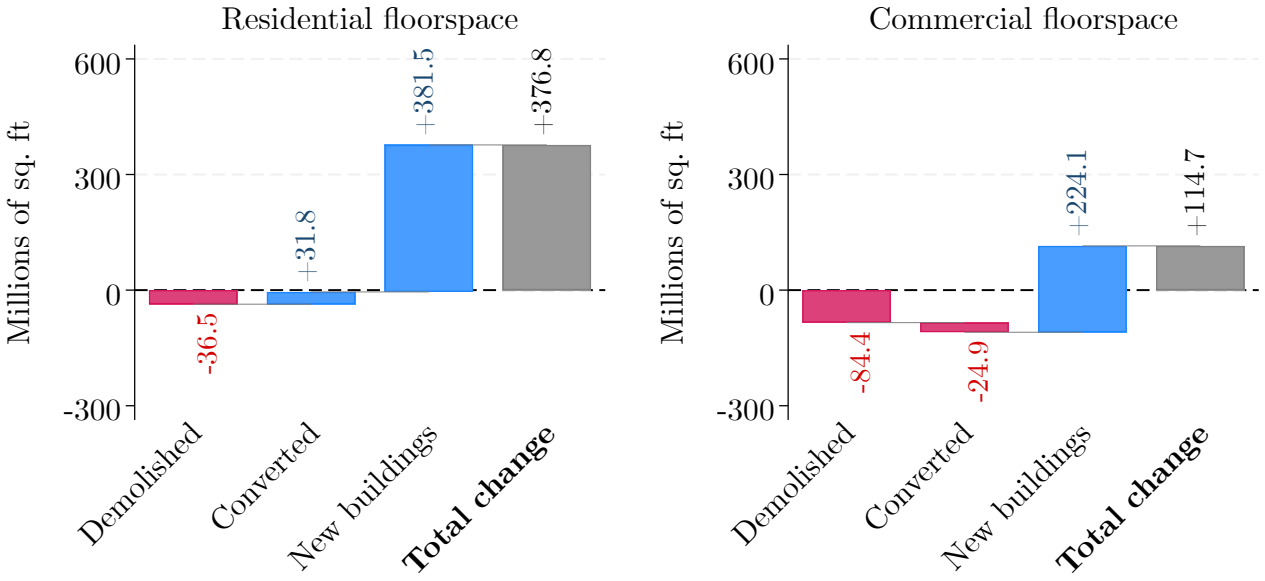
Figure B.8: Use restrictions likely hinder the reallocation of land uses in NYC



Notes: This figure illustrates the correlation between commercial-residential price differentials, plotted in Figure 4(b), and the share of land zoned for residential purposes in each neighborhood. Observations are weighted by the amount of floorspace in each neighborhood.

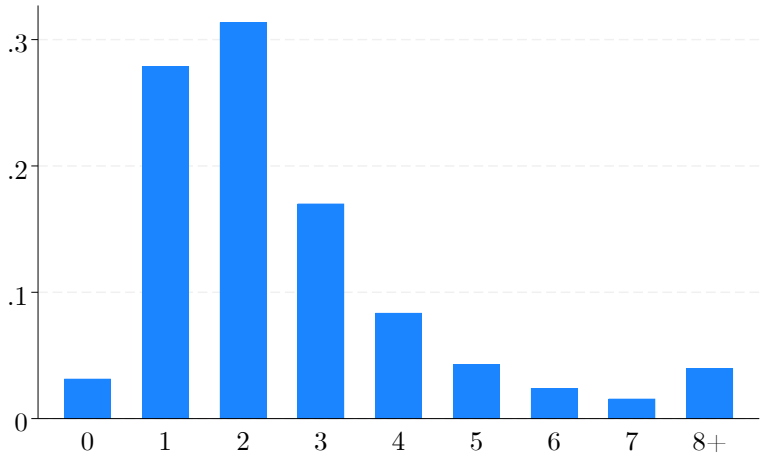
C Additional Descriptive Statistics

Figure C.1: Changes in floorspace, 2004–2022



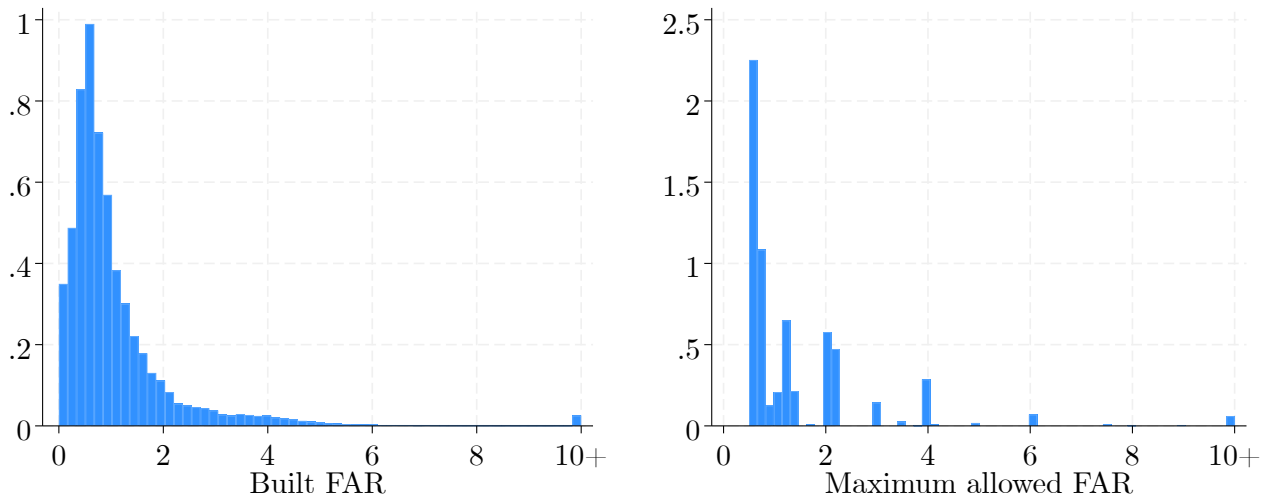
Notes: This figure decomposes the change in residential and commercial floorspace in NYC over 2004–2022 as the sum of demolished floorspace, converted floorspace, and newly built floorspace. In 2022, the city had 5.6 billion sq. ft of floorspace (1.9 billion sq. ft of commercial floorspace and 3.7 billion sq. ft of residential floorspace). The change in commercial floorspace caused by conversions is not exactly equal to minus the change in residential floorspace caused by conversions because the amount of floorspace in a building recorded in the tax data sometimes changes when conversions take place.

Figure C.2: Redevelopment duration (in years)



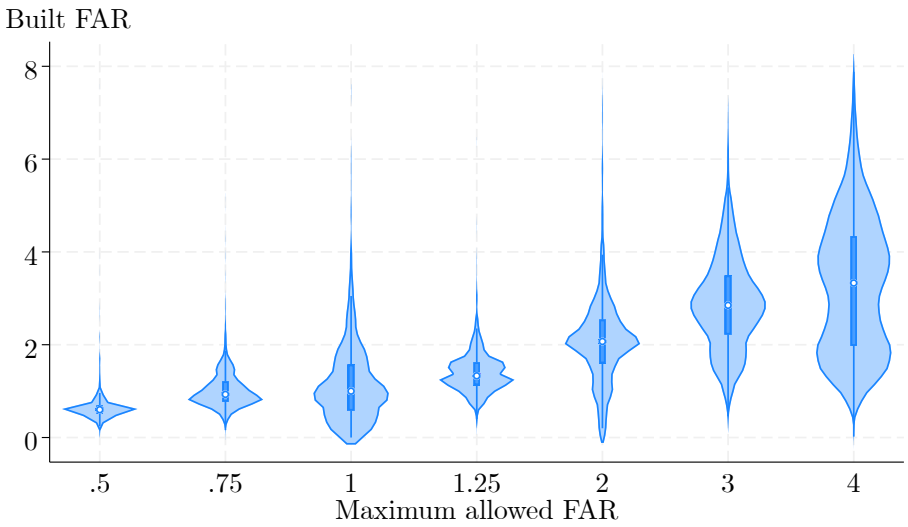
Notes: This figure plots the distribution of the duration of redevelopment projects (in years). I measure the duration of a project as the time elapsed between the approval of the building permit and the issuance of a certificate of occupancy, indicating that the building is suitable for use.

Figure C.3: Distribution of parcel FARs (built and allowed)



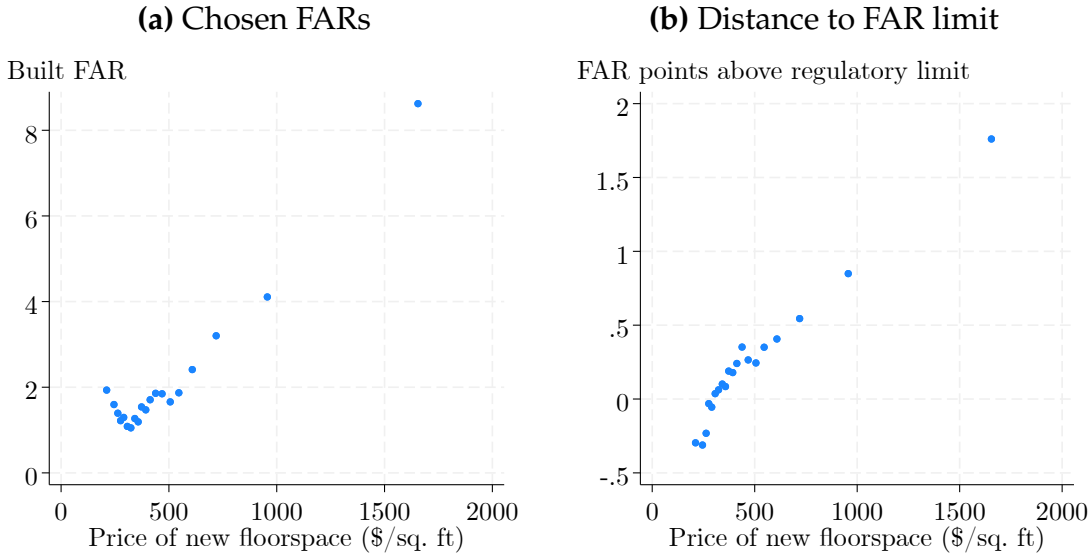
Notes: The left (resp., right) panel plots the distribution of built (resp., allowed) FARs in NYC in 2022.

Figure C.4: Built FARs and maximum allowed FARs



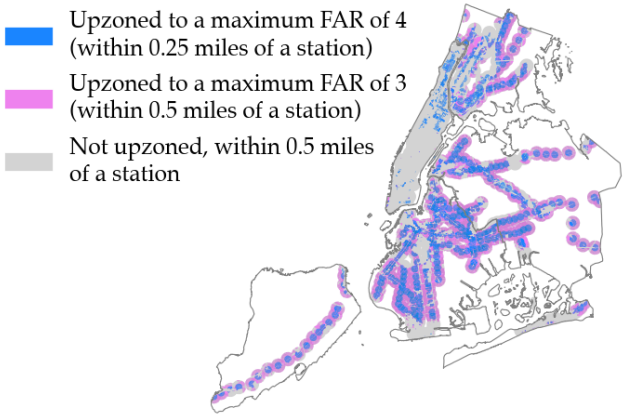
Notes: This figure plots the distribution the FAR of new buildings for seven common levels of maximum allowed FAR: 0.5, 0.75, 1, 1.25, 2, 3, and 4.

Figure C.5: Developers' choices and floorspace prices



Notes: This figure illustrates relationships between the choices made by developers and the floorspace prices they face. Panel (a) illustrates, for newly constructed buildings, the positive relationship between floorspace prices and the FARs they choose to build. Panel (b) shows the relationship between floorspace prices and the distance between the FAR chosen by developers and the maximum allowed FAR in the zoning code. Negative values correspond to developers building below the FAR limit. Developers often choose FAR levels slightly above the nominal FAR limit because some floorspace does not count towards the FAR limit (e.g., attics, cellars, or mechanical space). They can also obtain an additional FAR allowance if they invest in a local public good, buy "air rights", or request a slight relaxation of zoning rules (a zoning variance).

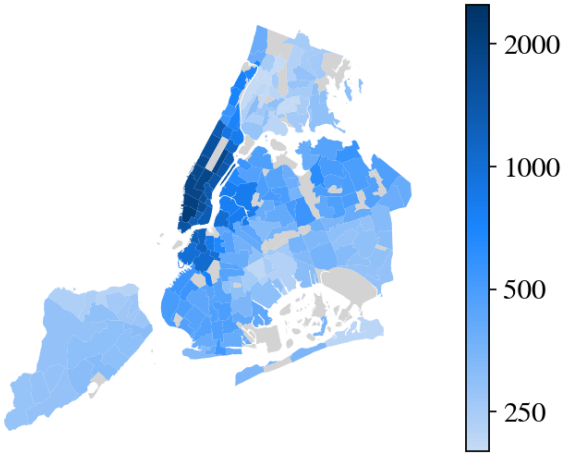
Figure C.6: Transit-Oriented Development counterfactual



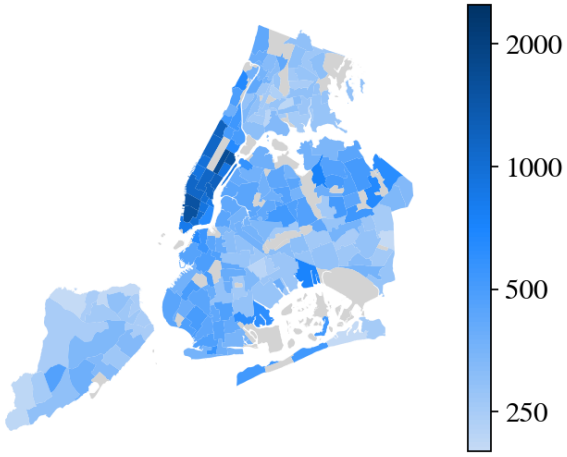
Notes: This figure maps the parcels upzoned in the Transit-Oriented Development counterfactual studied in Section 6. In this counterfactual, parcels within 0.25 miles of a subway, SIR, LIRR, Metro-North, or PATH station are upzoned to an allowance of 4 FAR points for both residential and commercial uses. Furthermore, the allowance for residential buildings is raised to 3 FAR points within 0.5 miles of a station. Parcels that have higher FAR allowances at baseline are not upzoned, leading most parcels in Manhattan to remain unaffected by the policy despite their proximity to transit stations. Furthermore, I exclude from the proposed upzoning parcels with landmarks, and those in historic districts or flood zones.

Figure C.7: Floorspace price levels ($/sq. ft)

(a) Residential prices



(b) Commercial prices



Notes: This figure shows prices of residential and commercial floorspace, adjusted for quality using the hedonic regression described in Section 5. Values shown correspond to average prices over 2004–2019.

Table C.1: Parcel characteristics

	Mean	P1	P10	P25	P50	P75	P90	P99
Lot area (sq. ft)	6,688	856	1,742	2,000	2,533	4,000	6,375	50,040
Floorspace (sq. ft)	6,458	0	1,056	1,377	2,000	3,032	5,933	91,115
Residential floorspace (sq. ft)	4,286	0	0	1,254	1,814	2,640	4,079	58,612
Commercial floorspace (sq. ft)	2,172	0	0	0	0	0	1,200	33,057
Building age (years)	80	7	33	64	89	99	109	139
Max. allowed FAR (residential uses)	1.3	0.0	0.5	0.5	0.8	2.0	3.0	7.5
Max. allowed FAR (commercial uses)	0.3	0.0	0.0	0.0	0.0	0.0	1.0	6.0
Max. allowed FAR	1.4	0.5	0.5	0.5	0.8	2.0	3.0	9.0
Parcel FAR	1.1	0.0	0.3	0.5	0.7	1.2	2.1	5.4
Binding FAR limit	0.612	0	0	0	1	1	1	1
Landmark or historical district	0.039	0	0	0	0	0	0	1
Redeveloped over 2004–2019	0.027	0	0	0	0	0	0	1

Notes: This table reports summary statistics for the parcels in my panel of land use in NYC. All variables are measured in 2019. I consider that the FAR limit is binding for a parcel if its FAR is above the maximum allowed FAR, or within 0.25 FAR points of the maximum allowed FAR.

D Supply of Floorspace: Additional Estimation Details

D.1 Variable costs estimation

Developer's problem. Developers choose the FAR of new buildings to maximize the variable profit from construction under the zoning constraint:

$$\max_{h^{\text{new}} \leq (1+b)\bar{h}^{\text{code}}_{\theta it}} P_{i,t}(h^{\text{new}}, \mathbf{x}^{\text{new}}_{it}) - VC^{\theta}_{it}(h^{\text{new}}_{it}) = p^{\theta}_{n,t}(h^{\text{new}})^{1+\beta^{\theta}_{\text{FAR}}} - \alpha^0 \alpha^{\theta}_b \alpha_t (h^{\text{new}})^{1/\zeta}, \tag{D.1}$$

Where $p^{\theta}_{nt} = \exp(\rho^{\theta}_{1,n} + \rho^{\theta}_{2,dt} + \sigma^2_\nu/2)$ is the price index for buildings of type θ in neighborhood n in year t, estimated in the first step of the estimation procedure, and $\beta^{\theta}_{\text{FAR}}$ is the hedonic coefficient on $\log(\text{FAR})$ in $\boldsymbol{\beta^{\theta}}$. As h^{new} increases, the marginal value of a sq. ft of floorspace decreases (when $\beta^{\theta}_{\text{FAR}}$ is negative), for instance because adding stories to a building requires allocating a larger share of the floorspace to elevators. The marginal construction cost of this floorspace also increases, and, in the absence of zoning constraints, the developer chooses an FAR h^* equating the marginal value of floorspace with its marginal cost:

$$h^* = \left[\frac{(1+\beta^{\theta}_{\text{FAR}})\zeta p^{\theta}_{n,t}}{\alpha^0 \alpha^{\theta}_b \alpha_t e^{\eta_{it}}} \right]^{\frac{1}{1/\zeta - 1 - \beta^{\theta}_{\text{FAR}}}}. \tag{D.2}$$

If $h^* > (1+b)\bar{h}^{\text{code}}_{\theta it}$, the zoning constraint will be binding and the developer will choose to build up to this limit.

Likelihood function. If we observe a developer choosing an FAR $h \leq h^{\text{code}}_{\theta it}$, we know that zoning was not a binding constraint, and that the developer drew a cost shock

$$\eta^* = \log\left(\frac{(1+\beta^{\theta}_{\text{FAR}})\zeta p^{\theta}_{n,t}}{\alpha^0 \alpha^{\theta}_b \alpha_t} \right) + \left(1 + \beta^{\theta}_{\text{FAR}} - \frac{1}{\zeta}\right) \log(h), \tag{D.3}$$

And the likelihood of this observation is

$$\ell = -\frac{1}{h}\left(1 + \beta^{\theta}_{\text{FAR}} - \frac{1}{\zeta}\right) \varphi_{\sigma_\eta}(\eta^*), \tag{D.4}$$

Where φ_{σ_η} is the PDF of a normal distribution with standard deviation σ_η. However, if $h > h^{\text{code}}_{\theta it}$, the FAR limit may have binded. In that case, the likelihood of observing a chosen FAR of h is given by

$$\ell = \left(-\frac{1}{h}\left(1 + \beta^{\theta}_{\text{FAR}} - \frac{1}{\zeta}\right) \varphi_{\sigma_\eta}(\eta^*) + \Phi_{\sigma_\eta}(\eta^*) \frac{\lambda}{h^{\text{code}}_{\theta it}} \right) \exp\left[-\lambda \left(\frac{h}{h^{\text{code}}_{\theta it}} - 1 \right) \right], \tag{D.5}$$

Where φ_{σ_η} is the CDF of a normal distribution with standard deviation σ_η.

D.2 Fixed costs estimation

Expected profits from redevelopment. To estimate fixed costs, I first compute for each parcel, year, and type of structure (residential or commercial) the developer's expected profit from redevelopment, net of fixed costs, $\tilde{\pi}$.

A developer receiving a profit shock η would choose, in the absence of zoning constraints, to build up to a FAR $h^*(\eta)$, given by equation (D.2). If $h^*(\eta) \leq \bar{h}^{\text{code}}$, then the developer will always choose a FAR $h^*(\eta)$ and will receive the following profit (supressing parcel, time, and building type indices for conciseness):

$$\tilde{\pi}(\eta) = p h^*(\eta)^{1+\beta_{\text{FAR}}} - P^{\text{old}} - \alpha h^*(\eta)^{1/\zeta} e^{\eta}, \tag{D.6}$$

Where $P^{\text{old}} = P_{it}(h_{it}^{\text{old}}, \mathbf{x}_{it}^{\text{old}})$ is the value of the structure that is redeveloped. If $h^*(\eta) > \bar{h}^{\text{code}}$, however, the developer will be constrained by zoning if the b shock it receives is low enough. In that case, the developer's expected profit is given by

$$\begin{aligned}\tilde{\pi}(\eta) = &\underbrace{\int_0^{\bar{b}(\eta)} \left[p(\bar{h}^{\text{code}}(1+b))^{1+\beta_{\text{FAR}}} - P^{\text{old}} - \alpha(\bar{h}^{\text{code}}(1+b))^{1/\zeta} \right] \lambda e^{-\lambda b} \mathrm{d}b}_{\text{When zoning is binding}} \\ &+ \underbrace{\int_{\bar{b}(\eta)}^{\infty} \left[p h^*(\eta)^{1+\beta_{\text{FAR}}} - P^{\text{old}} - \alpha h^*(\eta)^{1/\zeta} e^{\eta} \right] \lambda e^{-\lambda b} \mathrm{d}b}_{\text{When zoning is not binding}},\end{aligned} \tag{D.7}$$

Where $\bar{b}(\eta) = \frac{h^*(\eta)}{\bar{h}^{\text{code}}} - 1$ is the lowest draw of b that leaves the developer unconstrained by zoning. This expression can be more easily computed as

$$\begin{aligned}\tilde{\pi}(\eta) = &\left(\frac{p(\bar{h}^{\text{code}})^{1+\beta_{\text{FAR}}} e^{\lambda}}{\lambda^{1+\beta_{\text{FAR}}}} \left[\Gamma(2+\beta_{\text{FAR}}, \lambda) - \Gamma(2+\beta_{\text{FAR}}, (1+\bar{b}(\eta))\lambda) \right] \right. \\ &\left. - \alpha e^{\eta+\lambda} \left(\frac{\bar{h}^{\text{code}}}{\lambda} \right)^{1/\zeta} \left[\Gamma(1+\frac{1}{\zeta}, \lambda) - \Gamma(1+\frac{1}{\zeta}, (1+\bar{b}(\eta))\lambda) \right] \right) \\ &+ \left(p h^*(\eta)^{1+\beta_{\text{FAR}}} - \alpha h^*(\eta)^{1/\zeta} e^{\eta} \right) e^{-\lambda \bar{b}(\eta)} - P^{\text{old}}\end{aligned} \tag{D.8}$$

Where $\Gamma(\cdot, \cdot)$ is the incomplete gamma function. The expected profit of developers before drawing the cost shock η is then given by

$$\tilde{\pi} = \int \tilde{\pi}(\eta) \varphi_{\sigma_\eta}(\eta) \mathrm{d}\eta, \tag{D.9}$$

Which can be computed by numerical integration.

Value function. Equipped with a guess for the path of prices in each neighborhood (p_{nt}^{θ}), the path of density levels ($\bar{h}_{nt}$), and the fixed cost parameters (δ and σ), I can compute for each parcel and year the expected profit from redevelopment (ignoring the idiosyncratic profit shocks ϵ), $\pi_{it}^{\theta} = \tilde{\pi}_{it}^{\theta} - FC_{it}$, as well as the dispersion σ_{it} of developers' idiosyncratic shocks.

In each period, the option value of deciding not to redevelop and waiting for a better rede-

velopment opportunity in the future is given by

$$V_{it}^{\varnothing} = \frac{1}{1+r}\Bigg\{(1 - p_i^{\text{transaction}})V_{i,t+1}^{\varnothing} + p_i^{\text{transaction}}\sigma_{i,t+1}\log\left[\exp\left(\frac{V_{i,t+1}^{\varnothing}}{\sigma_{i,t+1}}\right) + \sum_{\theta\in\theta_{i,t+1}^{\text{allowed}}\subset\{R,C\}}\exp\left(\frac{\pi_{i,t+1}^{\theta}}{\sigma_{i,t+1}}\right)\right]\Bigg\}. \quad \text{(D.10)}$$

I compute these option values iteratively. I assume that at a sufficiently distant time horizon $\bar{T}$, $\pi_{i,t}^{\theta} = \pi_{i,\bar{T}}^{\theta}$ for all $t > \bar{T}$. Then, $V_{it}^{\varnothing}$ is constant for $t \geq \bar{T}$ and can be computed as a fixed point of equation (D.10). Given this terminal value of $V^{\varnothing}$, I can compute the full path of continuation values using equation (D.10) again. In estimation, I set $\bar{T}$ to 2100.

Fixed costs parameter estimation. The likelihood of observing the redevelopment of parcel i in year t into a structure of type θ is positive only if θ is an allowed use, and, in that case, given by

$$\ell_{it}^{\theta} = \frac{\exp\left(\pi_{it}^{\theta}/\sigma_{it}\right)}{\exp\left(V_{it}^{\varnothing}/\sigma_{it}\right) + \sum_{\theta\in\theta_{it}^{\text{allowed}}\subset\{R,C\}}\exp\left(\pi_{it}^{\theta}/\sigma_{it}\right)}. \quad \text{(D.11)}$$

To recover fixed cost parameters, I start with an initial guess for the full path of prices and density levels. I then estimate the δ and σ parameters using maximum likelihood. Using these parameters and the initial guess, I can predict the behavior of developers each period, compute the expected supply of floorspace over time in each neighborhood, use the demand model to estimate corresponding market clearing rents, and compute the associated path of floorspace prices. This provides a new initial guess which I can use to repeat this procedure until all parameters have converged.

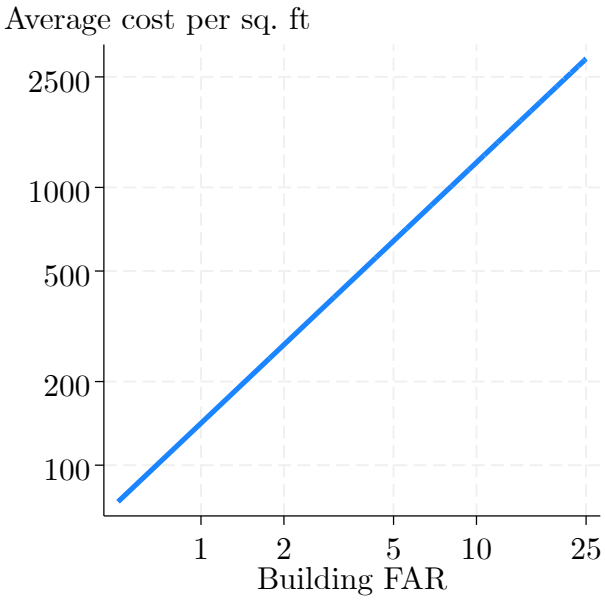
Multiple equilibria. The presence of agglomeration externalities raises the potential issue of multiple equilibria. For instance, if developers expect high future prices in a neighborhood, they will build more there, attracting more economic activity. With strong agglomeration forces, this can indeed lead to high prices. Conversely, under more pessimistic expectations, prices might remain at a lower equilibrium level.

To investigate the importance of multiple equilibria in my model, I follow Bayer and Timmins (2005) and Almagro et al. (2024) and compute equilibrium allocations by using a range of starting values for the solving algorithm. Specifically, I initialize the solver with extreme values across various neighborhoods, which increases the likelihood of discovering multiple equilibria. Using this procedure, I do not find evidence for meaningfully different multiple equilibria.

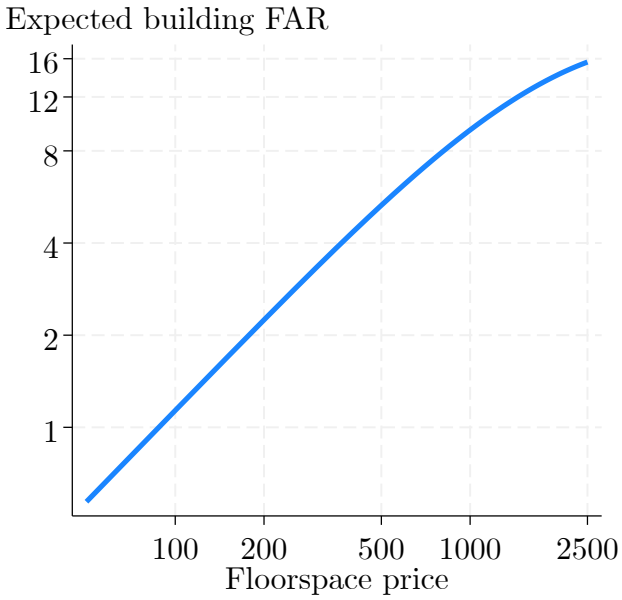
D.3 Parameter estimates and model fit

Figure D.1: Estimated variable costs

(a) Average costs per sq. ft for buildings of different FARs

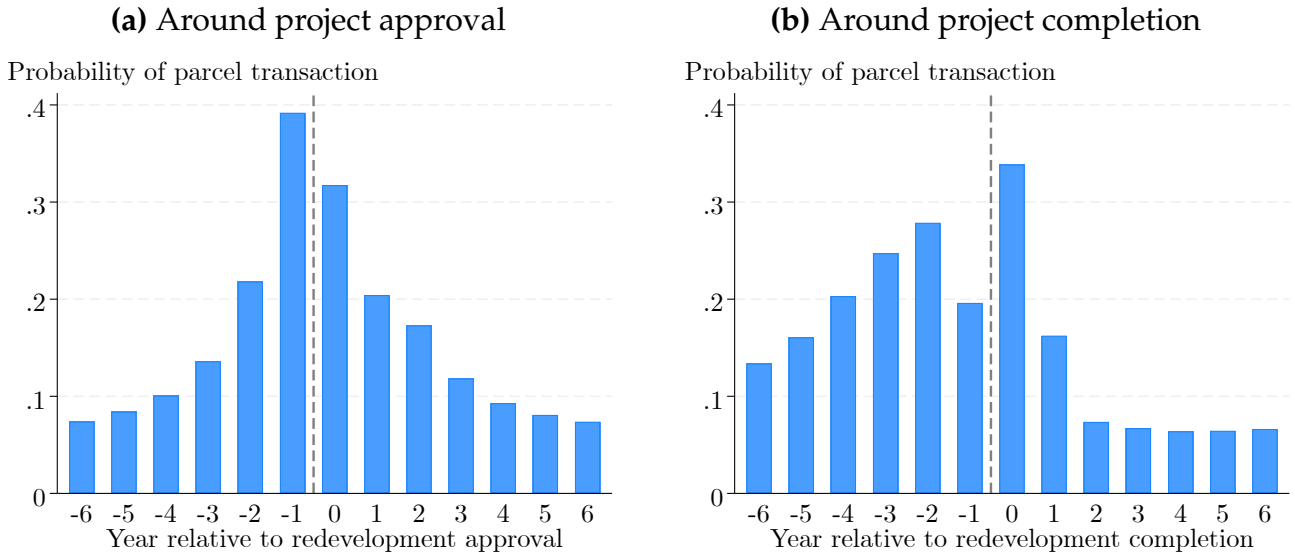


(b) Expected FAR of new buildings as a function of floorspace prices (without zoning)



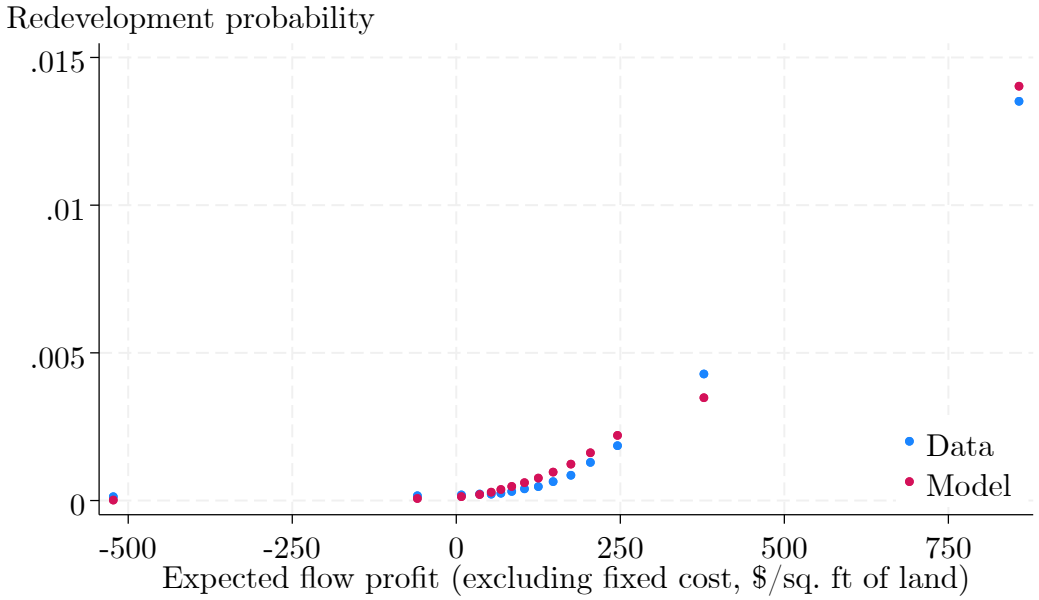
Notes: This figure presents estimation results for variable costs. Panel (a) shows the average construction cost that a developer has to pay to build a structure at different FAR levels, and panel (b) shows the expected FAR of a new building for different price levels faced by developers. In both panels, I use the estimated construction costs to build residential structures in Manhattan in 2019.

Figure D.2: Probability of parcel transaction around redevelopment events



Notes: This figure plots, for redeveloped parcels, the probability of observing these parcels being transacted in the sales data. We show probabilities for each year around the issuance of a building permit in panel (a) and for each year around the issuance of a certificate of occupancy in panel (b). Parcel transactions are measured in the sales data. I consider that a parcel is sold in a given year if any of its constituent zoning lots is sold that year, and I label a zoning lot as being sold during a year if more than 80% of its units were sold that year. The unconditional probability of a parcel being sold in a given year is 5.6%.

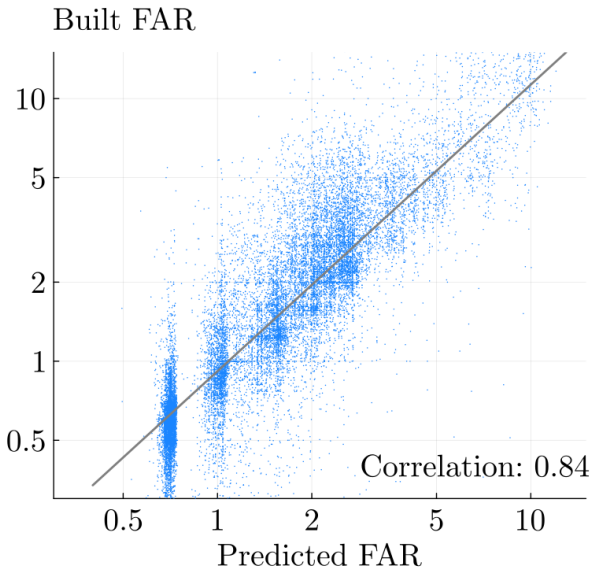
Figure D.3: Redevelopment probabilities and expected flow profits



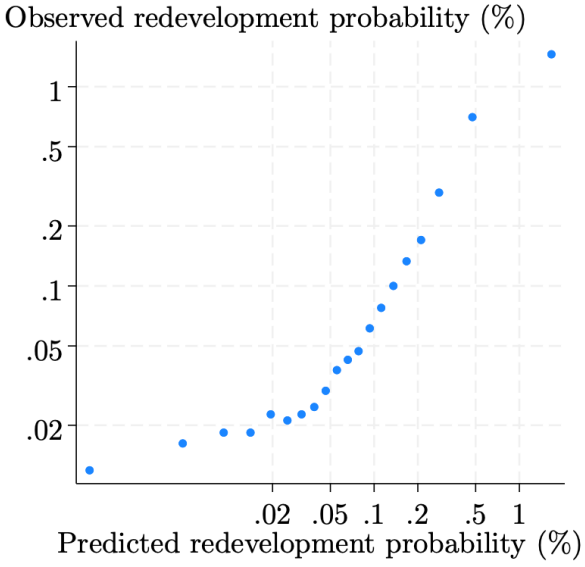
Notes: This figure shows the probability of redevelopment as a function of the expected flow profit from the most profitable redevelopment, $\max(\bar{\pi}^R_{it}, \bar{\pi}^C_{it})$, where $\bar{\pi}^\theta_{it} = \mathbb{E}\left[\max_{h^{\text{new}} \leq \bar{h}_{\theta it}} \left\{P_{it}(h^{\text{new}}, \mathbf{x}^{\text{new}}_{it}) - P_{it}(h^{\text{old}}_{it}, \mathbf{x}^{\text{old}}_{it}) - VC^\theta_{it}(h^{\text{new}}_{it})\right\}\right]$.

Figure D.4: Model fit, intensive and extensive margins

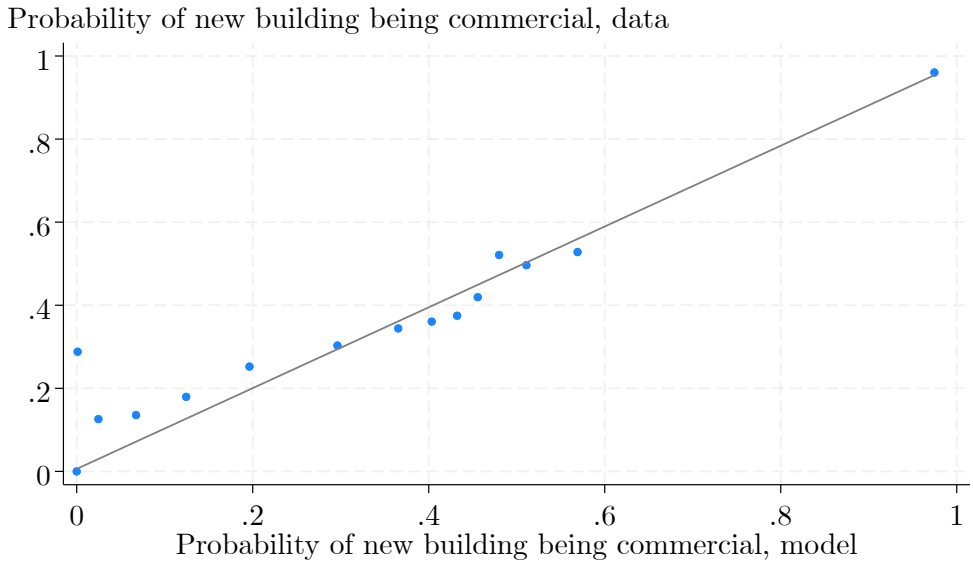
(a) Intensive margin: Predicted vs. observed FAR of new buildings



(b) Extensive margin: Predicted vs. observed redevelopment probabilities



(c) Probability of commercial vs. residential construction



Notes: Panel (a) compares, for buildings constructed between 2004 and 2019, the expected FAR of the structure according to the model with the observed built FAR. Each dot corresponds to a new structure, and the plotted line is the line of best fit. Panel (b) compares the model's predicted probability that a parcel will be redeveloped during a given year with the observed redevelopment probability in the data. Panel (c) compares the model's predicted probability that a developer chooses to build a commercial structure (conditional on redevelopment) with developers' observed choices.

Table D.1: Hedonic regression coefficients

	(1) Residential		(2) Commercial	
(log) Built FAR	-0.069	(0.002)	-0.098	(0.011)
(log) Unit size	-0.032	(0.002)	-0.107	(0.005)
Age	-0.002	(0.000)	0.000	(0.000)
Landmark	-0.133	(0.007)	0.083	(0.060)
Grade A	0.151	(0.003)	0.233	(0.032)
Grade B	-0.005	(0.002)	0.093	(0.019)
Grade C	-0.027	(0.002)	-0.024	(0.019)
Brick	-0.068	(0.003)	0.016	(0.083)
Frame	0.004	(0.003)	-0.226	(0.118)
Masonry	-0.246	(0.003)	-0.073	(0.018)
Office building			0.146	(0.025)
Retail building			0.217	(0.021)
Garage building			-0.007	(0.025)
Industrial building			0.022	(0.025)
Hotel			0.340	(0.048)
Neighborhood FE	Yes		Yes	
District $\times$ Year FE	Yes		Yes	
Observations	430,416		15,802	

Notes: This table reports estimation results for the hedonic model of equation (9). Unit size corresponds to the size of the unit sold if only one unit is sold and to the average unit size if several are. Brick, Frame, and Masonry are dummies indicating the building's main construction material.

Table D.2: Estimation of parcel sale probabilities

	Probit coefficients	
Parcel sold		
Office space (% of total floorspace)	-0.077	(0.006)
Retail space (% of total floorspace)	-0.018	(0.004)
Garage space (% of total floorspace)	-0.000	(0.007)
Storage space (% of total floorspace)	0.033	(0.009)
Factory space (% of total floorspace)	0.001	(0.007)
Hotel space (% of total floorspace)	-0.015	(0.020)
Other space (% of total floorspace)	-0.327	(0.006)
Condo/Coop	-0.684	(0.008)
Parcel in Bronx	0.051	(0.003)
Parcel in Brooklyn	0.038	(0.003)
Parcel in Queens	0.036	(0.003)
Parcel in Staten Island	0.023	(0.003)
Constant	-1.615	(0.003)
Observations	13,158,512	

Notes: This table reports estimated coefficients of a probit regression of a dummy indicating whether a parcel was sold in a given year on parcel characteristics.

E Demand for Floorspace: Additional Estimation Details

E.1 Benchmark model

In this section, I calibrate a simple quantitative spatial model where agents have homogenous skills and homothetic preferences, and where floorspace quantities are a structural residual, as in Ahlfeldt et al. (2015). These structural residuals poorly match the floorspace quantities I observe in the data, suggesting that this benchmark model fails to correctly predict household demand for floorspace in the context of NYC.

Consider a simplified version of the model presented in Section 4.2 where all agents can provide a single unit of labor and have the following utility function:

$$U_{ijo} = \frac{B_i z_{ijo}}{d_{ij}} C_{ijo}^{1-\beta} H_{ijo}^{\beta}, \tag{E.1}$$

Where z_{ijo} is an idiosyncratic shock, drawn from Fréchet distribution with shape $\varepsilon = 4.4$ (I estimate this parameter following Dingel and Tintelnot, 2020, using data on commuting flows and commuting times), and all agents spend a share $\beta = 0.33$ (calibrated from the ACS) of their income on housing. Other variables are defined as in Section 4.2.

Using data on the distribution of residents and jobs, I can calibrate the wages offered in each location as in Ahlfeldt et al. (2015) as the solution to the following system of equations:

$$L_{Fj} = \sum_i \left(\frac{(w_j / d_{ij})^{\varepsilon}}{\sum_s (w_s / d_{is})^{\varepsilon}} L_{Ri} \right), \tag{E.2}$$

Where L_{Fj} denotes the number of jobs in location j and L_{Ri} is the number of residents of location i.

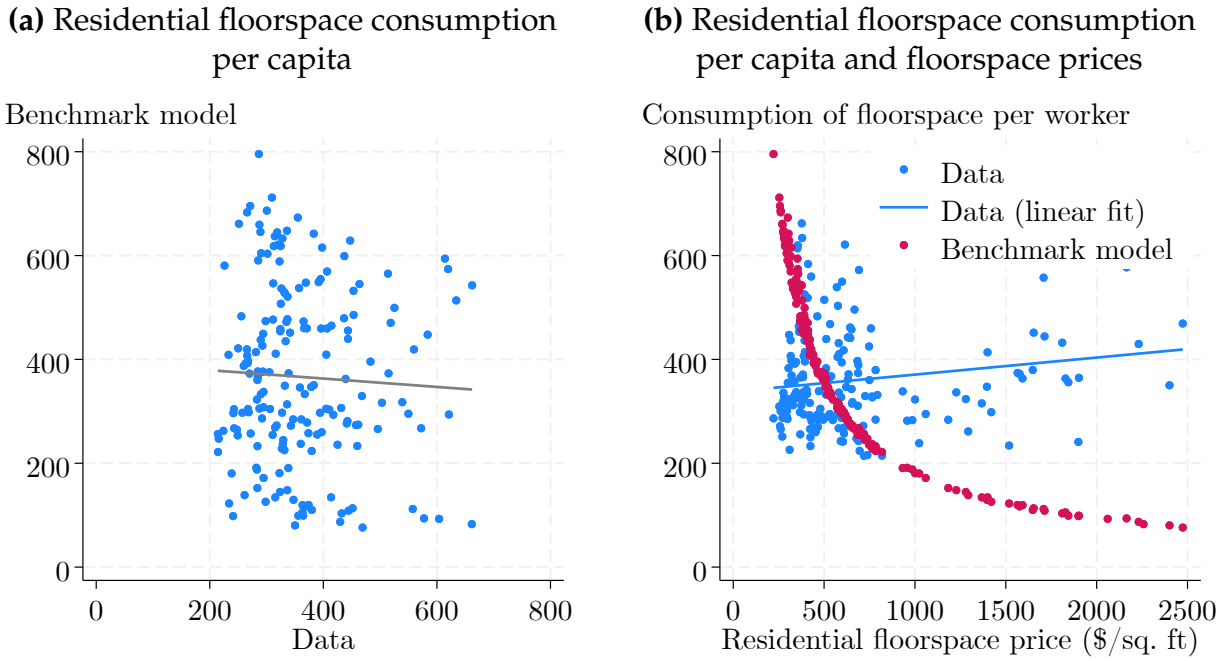
Individual incomes are only identifiable up to a constant, which I can calibrate using the total spending on housing in NYC, equal to β times the total income of residents of the city. I can then identify the total income of residents in each location, and given the prices $\boldsymbol{r_R}$ of floorspace, I can recover the average amount of floorspace consumed by person in each location $\tilde{\boldsymbol{H}}$.

In Appendix Figure E.1, I compare this structural residual with $\boldsymbol{H}$, the average amount of floorspace per capita observed in the data. There is little correlation between $\tilde{\boldsymbol{H}}$ and $\boldsymbol{H}$—if anything, the two variables are negatively correlated. Furthermore, the structural residual $\tilde{\boldsymbol{H}}$ is much more dispersed than the floorspace consumptions observed in the data. In the upscale SoHo-Tribeca neighborhood of Manhattan, the benchmark model predicts a housing consumption of 76 sq. ft per capita, about six times lower than the average consumption of 469 sq. ft observed in the data. Conversely, in the Hunts Point neighborhood in the Bronx, where about a third of the residents live below the poverty line, the benchmark model predicts an average floorspace consumption of 696 sq. ft, more than twice the 271 sq. ft observed in the data.

These examples, along with Appendix Figure E.1(b), illustrate why, in the context of NYC, the benchmark model fails to accurately predict floorspace consumption. In Appendix Figure E.1(b), I plot $\boldsymbol{H}$ and $\tilde{\boldsymbol{H}}$ as a function of residential floorspace prices $\boldsymbol{r_R}$. In the benchmark model, despite locations' heterogeneous access to high-wage workplaces, the average income of residents is relatively homogeneous across locations. Hence, agents' consumption of floorspace is approximately inversely proportional to residential prices. Because there is wide variation in prices across neighborhoods, agents' consumption of floorspace is also very heterogeneous in the benchmark model.

In reality, per capita housing consumption is relatively homogeneous across neighborhoods

Figure E.1: Benchmark model fit



Notes: This figure evaluates model fit for the benchmark model presented in Appendix Section E.1, where agents have homogeneous skill levels and homothetic preferences. In this model, the amount of residential floorspace per capita in each neighborhood is a structural residual. In panel (a), I compare this structural residual with the amount of residential floorspace observed in the data. In panel (b), I plot both variables against the price of residential floorspace.

and, if anything, tends to increase with residential prices. These patterns can be rationalized by strongly heterogeneous skill levels and housing being a necessity.[27] This will lead the rich to locate in expensive, high-amenity neighborhoods while poorer households households will sort to cheaper, low-amenity neighborhoods. These assumptions allow the model of Section 4.2 to match the housing consumption levels I observe in the data.

E.2 Additional model results

In this section, I derive a set of theoretical implications of the model described in Section 4.2 that allow its calibration.

Workers' choices. Given the timing of the idiosyncratic shocks z_i^H and z_j^W, I can solve for workers' decisions through backward induction. Consider a worker of type θ who chose home location i. Upon drawing shock z_j^W, they choose a work location j to maximize their utility. Choosing

[27] A model with two types of workers, e.g., college-educated vs. non-college-educated (as in Tsivanidis, 2019) fails to match the observed consumption patterns in NYC. Indeed, residents in the most expensive neighborhoods of NYC spend, on average, almost ten times more on housing than those in the cheapest neighborhoods of the city. A two-type model cannot satisfyingly explain these wide differences across neighborhoods. This is why I introduce a wider range of types, which captures the large heterogeneity in New Yorkers' earnings capacity.

to work in location j will yield them an income $I_{ij\theta} = w_j s(\theta)(1+t)$, a utility level

$$U_{ij\theta} \propto B_i z_i^H z_j^W u_{ij\theta} = B_i z_i^H z_j^W \max\left\{ \frac{I_{ij\theta} - r_{Ri}\underline{H}_i}{d_{ij} r_{Ri}^{\beta}}, 0 \right\}, \tag{E.3}$$

And a housing spending

$$r_{Ri} H_{ijo} = r_{Ri}\underline{H}_i + \beta(I_{ij\theta} - r_{Ri}\underline{H}_i). \tag{E.4}$$

Workers will only commute to locations with a wage high enough to provide them with an income larger than $r_{Ri}\underline{H}_i$, and within that set of locations, the probability that they choose location j is given by:

$$p_{j|i,\theta} = \frac{u_{ij\theta}^{\varepsilon^W}}{\sum_s u_{is\theta}^{\varepsilon^W}}. \tag{E.5}$$

A worker of type θ choosing to live in location i receives the following expected utility level before drawing z_j^W:

$$U_{i\theta} \propto B_i z_i^H u_{i\theta} = B_i z_{io}^H \left(\sum_j u_{ij\theta}^{\varepsilon^W} \right)^{1/\varepsilon^W} \tag{E.6}$$

If a worker of type θ cannot afford $r_{Ri}\bar{h}_i$ no matter its chosen work location j, their expected utility $U_{i\theta}$ in that location will be equal to zero. The probability that a worker of type θ will choose to live in location i is given by:

$$p_{i|\theta} = \frac{(B_i u_{i\theta})^{\varepsilon^H}}{\sum_s (B_i u_{i\theta})^{\varepsilon^H}}. \tag{E.7}$$

Before drawing their idiosyncratic shocks, the expected utility of a worker of type θ is given by

$$U_\theta = \left(\sum_i \left(B_i \left(\sum_j u_{ij\theta}^{\varepsilon^W} \right)^{1/\varepsilon^W} \right)^{\varepsilon^H} \right)^{1/\varepsilon^H}. \tag{E.8}$$

The average expected utility of a worker in the city is denoted $U = \int_\theta U_\theta \, \mathrm{d}F(\theta)$. Changes in the city's population are governed by the migration elasticity $\varepsilon_M = \mathrm{d}\log(L)/\mathrm{d}\log(U)$.

E.3 Calibration procedure

Worker types. To calibrate the distribution of skill levels, I rely on the ACS. I assume that $s(\theta)$ is distributed as a Lognormal$(0, 1.2)$ truncated at the 10th percentile. This allows me to match the observed income distribution in the ACS after excluding households with an annual income below \$12,000. To facilitate computation, I approximate this continuous distribution of skill levels with 100 discrete types.

I also rely on the ACS to calibrate β. I find that the richest households in the NYC metropolitan area spend about 10% of their income on housing, and therefore calibrate β to 0.1.

Preferences for work locations. In the absence of subsistence levels for housing $\underline{H}_i$, the probability that a worker living in i will choose to work in j is

$$p_{j|i} = \frac{(w_j/d_{ij})^{\varepsilon^W}}{\sum_s (w_s/d_{is})^{\varepsilon^W}}, \tag{E.9}$$

And the share of workers π_{ij} commuting between i and j is related to travel times between these locations through the following gravity equation, as in Ahlfeldt et al. (2015):

$$\log \pi_{ij} = -\nu \tau_{ij} + \vartheta_i + \varsigma_j, \tag{E.10}$$

Where $\nu = \kappa \varepsilon^W$. This equation can be estimated using data on commuting flows from LODES and commuting times from Google Maps. To account for the sparsity of the commuting flows matrix, I follow Dingel and Tintelnot (2020) and estimate equation (E.10) through PPML. I find a commuting elasticity ν approximately equal to 0.044. Assuming that the commuting cost parameter κ is equal to 0.01 as in Ahlfeldt et al. (2015), I set ε^W to 4.4.

The presence of subsistence levels $\underline{H}_i$ distorts agents' choices of work locations and makes equation (E.10) no longer hold. However, this distortion induced by non-homothetic preferences is minimal. After calibrating ε^W to 4.4 and other parameters using the procedure laid out below, I can predict commuting flows $\tilde{\pi}_{ij}$ between locations. Estimating equation (E.10) using $\tilde{\pi}_{ij}$ instead of π_{ij} yields an estimated coefficient of 0.0441, instead of the 0.044 that would be estimated when the subsistence levels $\underline{H}_i$ are set to zero. This justifies estimating ν using equation (E.10) despite it not strictly holding in the model.

Preferences for home locations. Then, I calibrate the dispersion of idiosyncratic preferences for home locations ε^H, the subsistence levels of housing $\underline{\boldsymbol{H}}$, amenity levels $\boldsymbol{B}$, and wages $\boldsymbol{w}$. To facilitate estimation, I introduce $\tilde{\boldsymbol{w}} = \boldsymbol{w} / \left(\prod w_j\right)^{1/S}$, a normalized vector of wages such that the geometric mean of $\tilde{\boldsymbol{w}}$ is one, and an income normalization parameter $\mathcal{I}$ such that the income of a worker of type θ working in location j is $I_{ij\theta} = \mathcal{I} \cdot s(\theta) \tilde{w}_j$. To calibrate these parameters, I use the following iterative procedure:

1. I start with an initial guess of parameters $(\varepsilon^H, \underline{\boldsymbol{H}}, \boldsymbol{B}, \tilde{\boldsymbol{w}}, \mathcal{I})$, as well as for the distribution of types among the residents of each location.

2. The guessed parameters imply a distribution of work locations through equation (E.5). Increasing the wage in a location increases the number of people working in that location. I update $\tilde{\boldsymbol{w}}$ until I match the number of jobs in each location reported by LODES, keeping the geometric mean of $\tilde{\boldsymbol{w}}$ normalized to one.

3. The guessed parameters imply a distribution of home locations through equation (E.7). Increasing the amenity level in a location increases the number of people living there. I update $\boldsymbol{B}$ until I match the number of residents in each location reported by LODES. Because $\boldsymbol{B}$ is only identified up to a constant, I normalize its geometric mean to one.

4. The guessed parameters imply a total spending on housing in each neighborhood through equation (E.4). Increasing the subsistence housing level $\bar{h}$ in a location increases the housing spending in that location. I update $\underline{\boldsymbol{H}}$ until the total spending on housing in each neighborhood $\boldsymbol{r_R} \circ \boldsymbol{H_R}$ in the data is matched. At each update of $\underline{\boldsymbol{H}}$, I update the guess of

the distribution of types in each location using equation (E.7), as well as the guess of wages and amenities by going through steps (2) and (3).

5. The guessed parameters imply a distribution of the share of income spent on housing $r_R H/I$. Increasing the income normalization parameter $\mathcal{I}$ decreases the average share of income spent on housing. I update $\mathcal{I}$ until the average worker spends 33% of their income on housing, which corresponds to the average spending on housing observed in the ACS. At each update of $\mathcal{I}$, I iterate through steps (2)–(4) to concurrently update wages, amenities, and subsistence levels.

6. The guessed parameters imply a distribution of the average income level in each neighborhood. Increasing ε^H decreases the heterogeneity in workers' idiosyncratic preferences for home locations, leads to a stronger sorting of high-skill workers to high-amenity locations, and increases the dispersion of neighborhoods' average income. I update ε^H until the standard deviation of (log) average neighborhood per capita income matches that observed in the ACS (of 0.51). For each update of ε^H, I iterate through steps (2)–(5) to concurrently update wages, amenities, subsistence levels, and the income normalization parameter.

7. Once the parameters $\varepsilon^H, \underline{\boldsymbol{H}}, \boldsymbol{B}, \tilde{\boldsymbol{w}}$, and $\mathcal{I}$ are calibrated, I can compute the total income of workers I_{total}. That income is the sum of total wage income I_{wages} and total income from rents, $I_{\text{rents}} = \sum_i (r_{Ri} H_{Ri} + r_{Fi} H_{Fi})$ which is observable in the data. With a measure of I_{total} and I_{rents}, I can recover $t = I_{\text{rents}}/I_{\text{wages}} = I_{\text{rents}}/(I_{\text{total}} - I_{\text{rents}})$. Knowing t then allows me to recover the vector of wages, as $w_j = (\mathcal{I}\tilde{w}_j)/(1+t)$.

Productivities and floorspace shares. Once wages and other parameters describing workers' preferences have been calibrated, I can compute the effective labor supply $\boldsymbol{L_F}$ provided in each location and infer the vector $\boldsymbol{\alpha}$ through the share of floorspace in firms' costs in each location:

$$\alpha_j = \frac{R_{Fj} H_{Fj}}{R_{Fj} H_{Fj} + w_j L_{Fj}}. \tag{E.11}$$

Finally, the productivity in each location can be inferred as

$$A_j = \frac{w_j}{1-\alpha_j} \left(\frac{L_{Fj}}{H_{Fj}} \right)^{\alpha_j}. \tag{E.12}$$

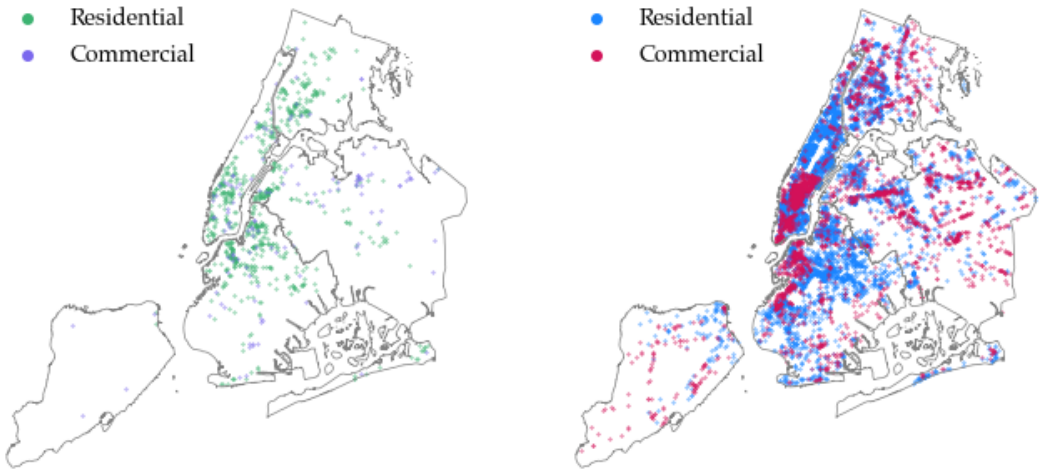
Outside location. I estimate the demand model using NYC's partition into 188 Neighborhood Tabulation Areas (NTAs), corresponding to historical neighborhoods of the city (e.g., Chinatown, Hamilton Heights). NYC defines seven additional NTAs that correspond to parks, cemeteries, airports, and Rikers Island—these neighborhoods and the parcels within them are excluded from estimation. I add to the model an outside location corresponding to the rest of the metropolitan area, which workers in the model can commute to and from. LODES indicates the number of workers living and working in this additional location, but I do not have data on the supply of floorspace there as well as its price. I set the initial amount of residential (resp., commercial) floorspace in the outside location such that the amount of floorspace per worker (resp., per job) is at the same level as in NYC's outer boroughs. Furthermore, I set the price of floorspace in this outside location to the average price of residential floorspace in the city's outer boroughs. In

counterfactual simulations, I assume the supply of floorspace in the outer location grows at the same rate as in the city.

E.4 Estimation of spillover parameters

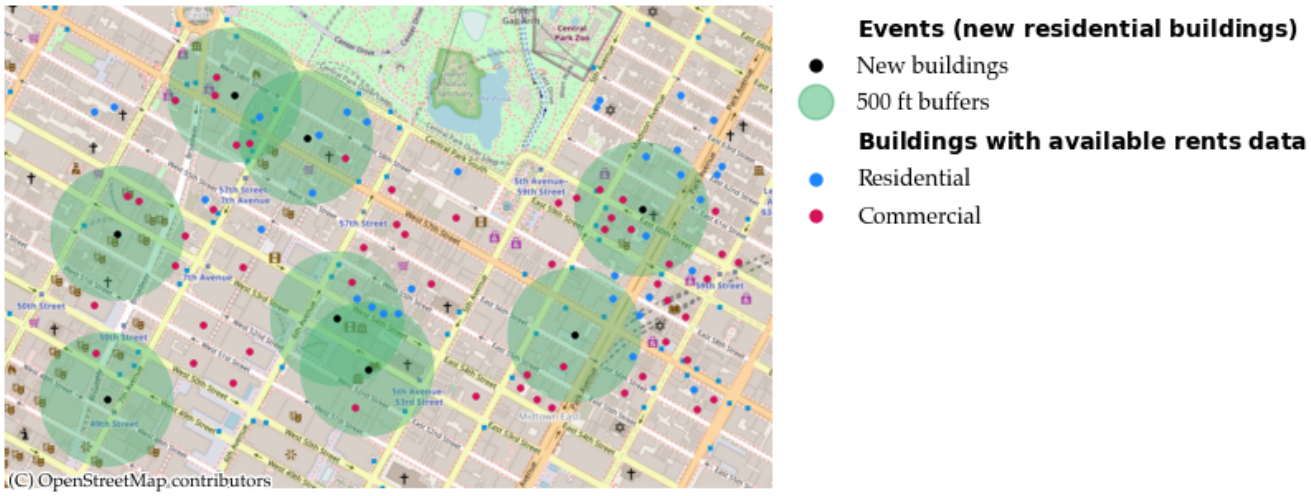
Figure E.2: Location of large construction events and buildings with rent data

(a) Construction events **(b)** Buildings with rent data



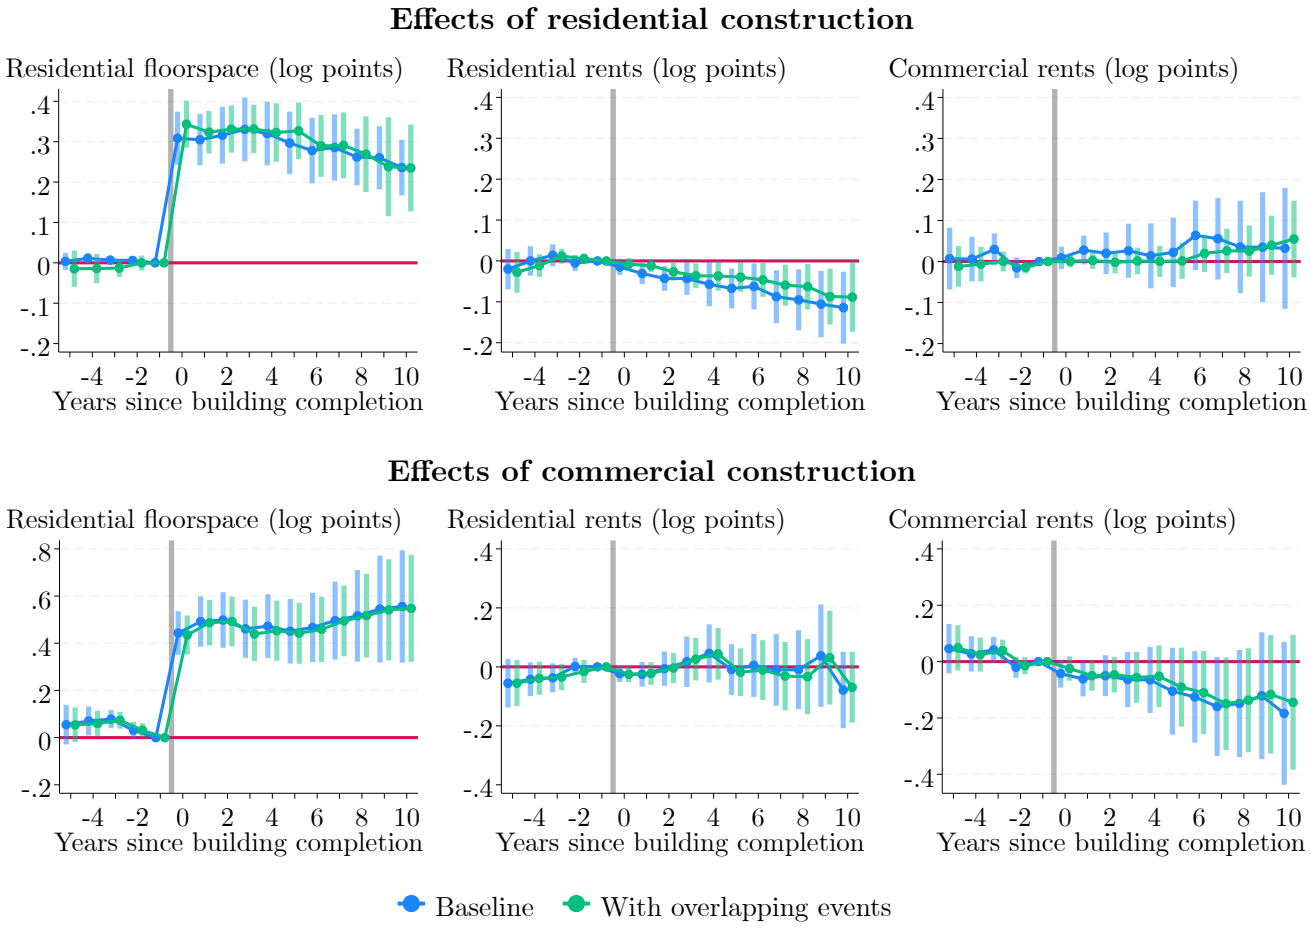
Notes: The left panel shows the locations of the new construction events leveraged in the event studies of Figure 6. The right panel shows the locations of buildings for which rent data is available over 2004–2021 and that are not rent-stabilized.

Figure E.3: Examples of residential construction events and associated buffers



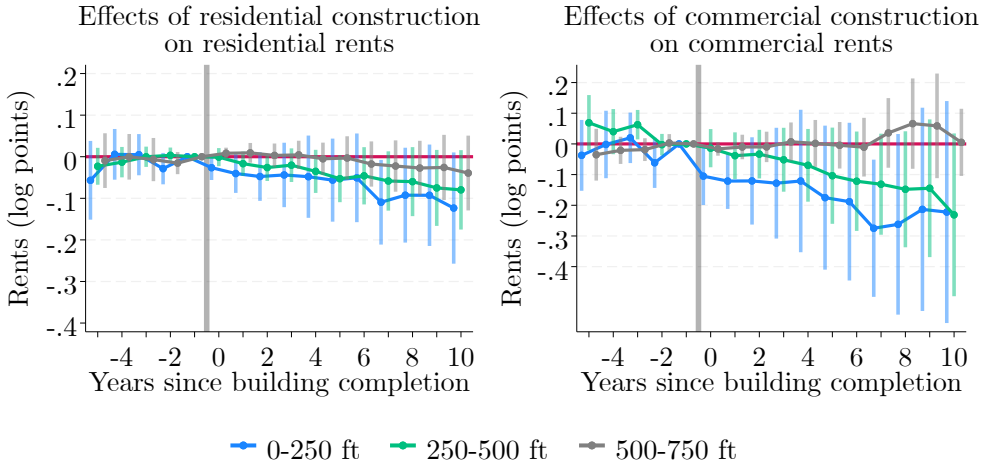
Notes: This figure illustrates the event studies of Figure 6 in a section of Midtown Manhattan. The black dots correspond to the large residential construction projects used as events. The green disks correspond to 500 ft buffers around new buildings, in which we measure rent levels for existing buildings. The blue (resp., red) dots correspond to residential (resp., commercial) buildings for which we have data on rents from NOPVs. Rent-stabilized buildings are excluded. In dense areas of the city like the one pictured here, some of the 500 ft buffers used in event studies overlap. Appendix Figure E.4 shows the sensitivity of the event study results to including these overlapping events.

Figure E.4: Robustness to including overlapping buffers



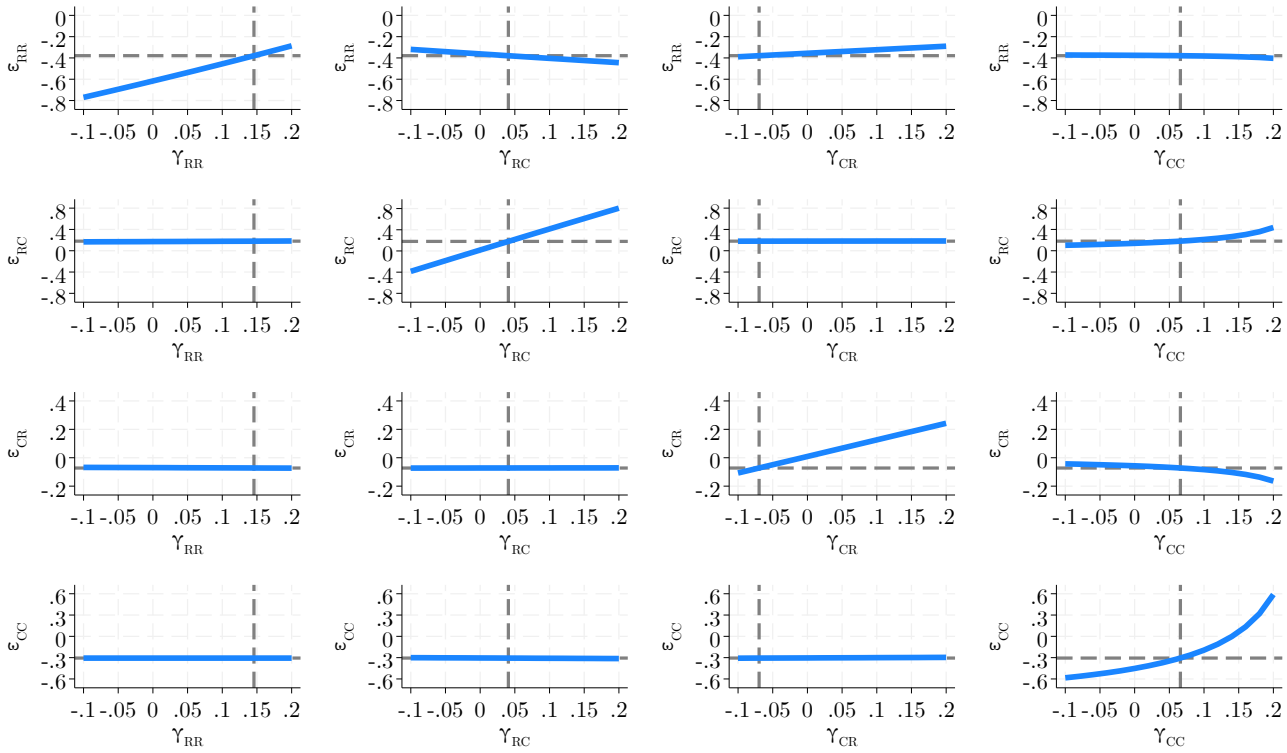
Notes: This figure shows the estimates of Figure 6 (in blue) as well as analog estimates when including from the set of new building events those located within 500 ft of another event (in green).

Figure E.5: Spatial decay of price effects



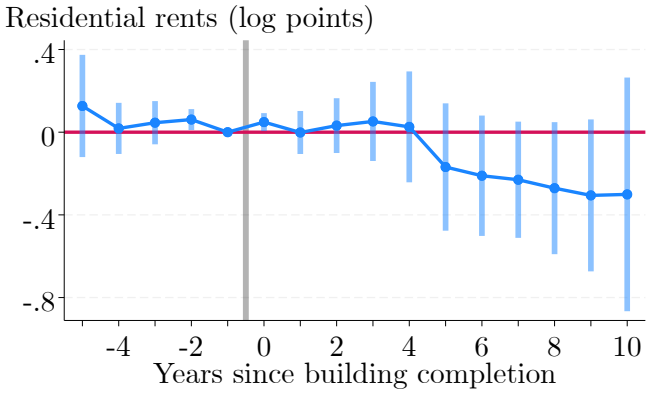
Notes: The left (resp., right) panel shows the effects of new residential (resp., commercial) construction on residential (resp., commercial) rents, measured at different distances from the studied event. In blue, I show effects on rents measured within 250 ft of the new building; in green, I show effects on rents measured between 250 and 500 ft of the new building; and in gray, I show effects on rents measured between 500 and 750 ft of the new building.

Figure E.6: Sensitivity of reduced-form elasticities to spillover parameters



Notes: This figure shows how, in the model, the reduced-form parameters ε_{RR}, ε_{RC}, ε_{CR}, and ε_{CC}, describing the elasticity of residential/commercial rents in a neighborhood to the supply of residential/commercial floorspace there, vary with the spillover parameters γ_{RR}, γ_{RC}, γ_{CR}, and γ_{CC}, which determine how amenity/productivity levels in a neighborhood vary with the density of workers/jobs there. In each panel, I show how an ε parameter varies with the values of a γ parameter, keeping the values of the other spillover parameters to their baseline value (indicated by the vertical dashed lines). The values of the ε parameters estimated in Figure 6 are denoted by the horizontal dashed lines.

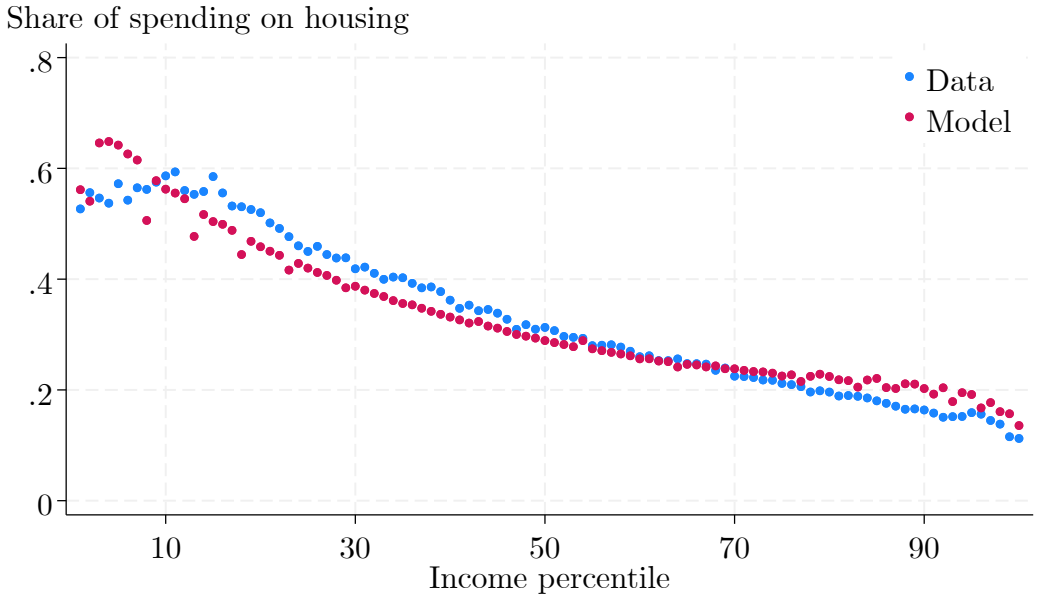
Figure E.7: Effects of industrial construction on residential rents



Notes: This figure shows results for an exercise similar to that of Figure 6, estimating the effect of building new large industrial buildings (factories, garages, warehouses) on the rents of nearby residential buildings.

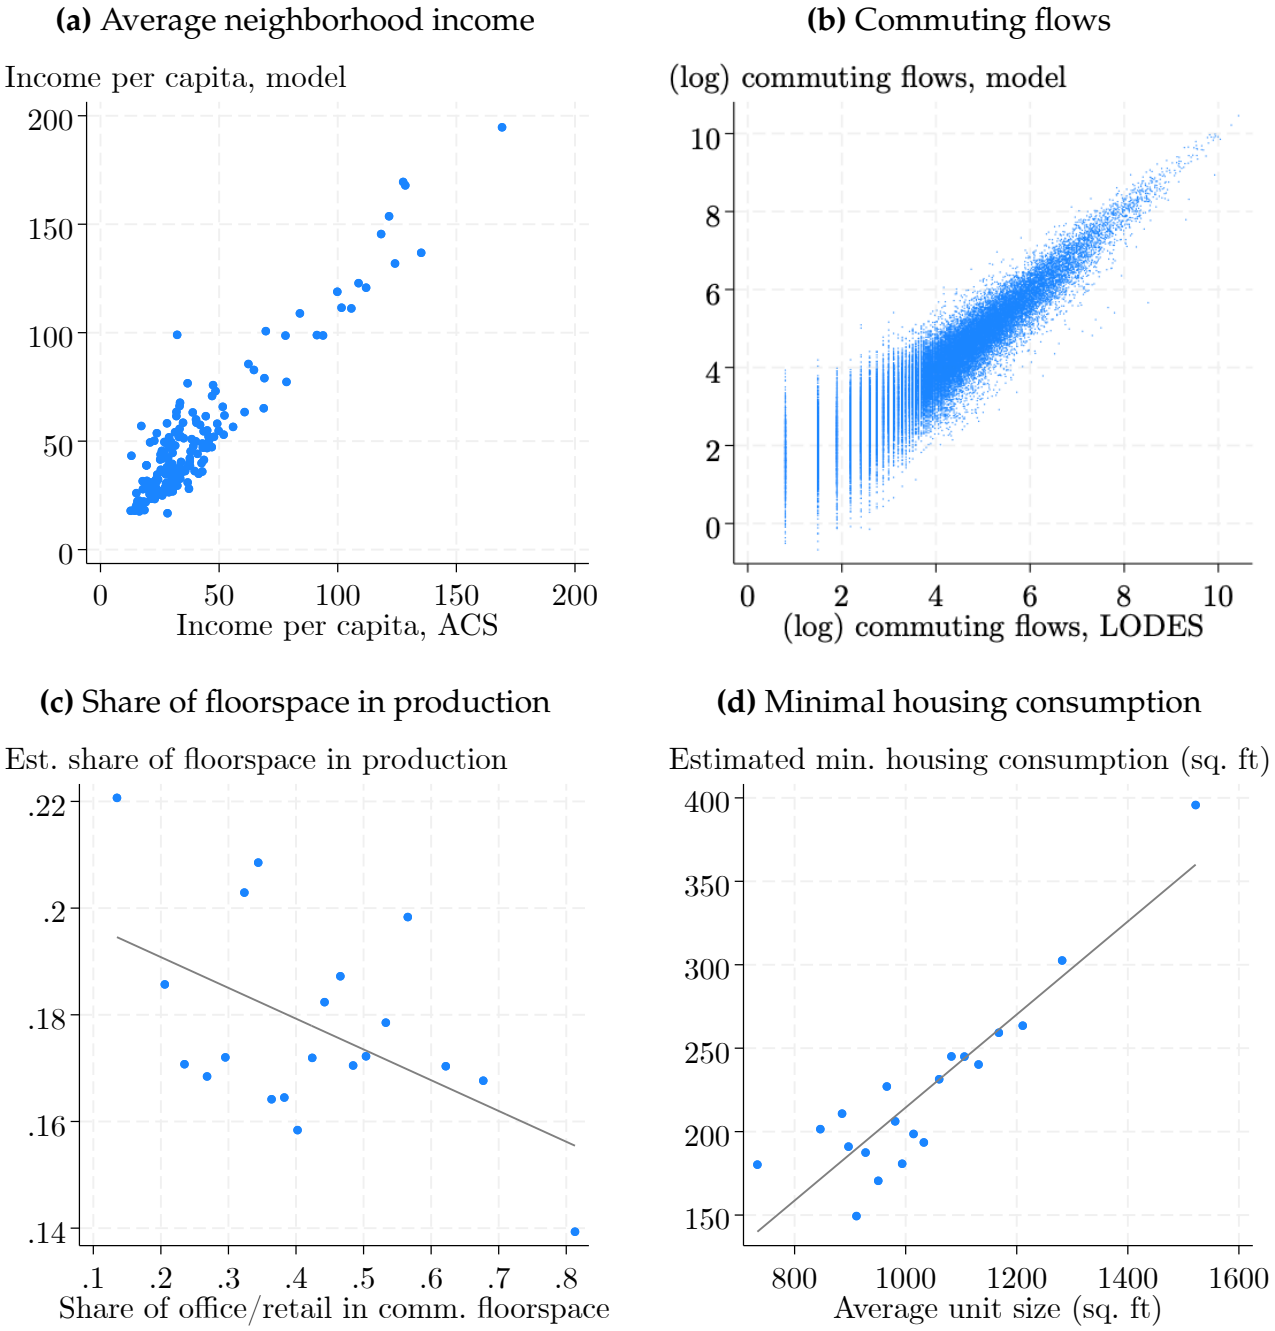
E.5 Parameter estimates and model fit

Figure E.8: Engel curve for housing



Notes: This figure shows the share of income spent on housing at different percentiles of the income distribution in the model and the NYC metropolitan area sample of the 2019 5-year ACS.

Figure E.9: Demand model fit



Notes: Panel (a) compares commuting flows between each neighborhood in LODES and in the model. Panel (b) compares average neighborhood income in the data and in the ACS. Panel (c) compares the estimated share of floorspace in production in each neighborhood (α in the model) and the share of commercial floorspace that is used for office or retail uses (as opposed to garage, storage, factory, or other commercial uses), reported in the PLUTO dataset. Panel (d) compares the estimated minimal amount of housing consumption in each neighborhood ($\bar{h}$ in the model) with the average size of residential units reported in the PLUTO dataset.

Table E.1: Correlation between model-derived amenities and survey-reported neighborhood satisfaction levels

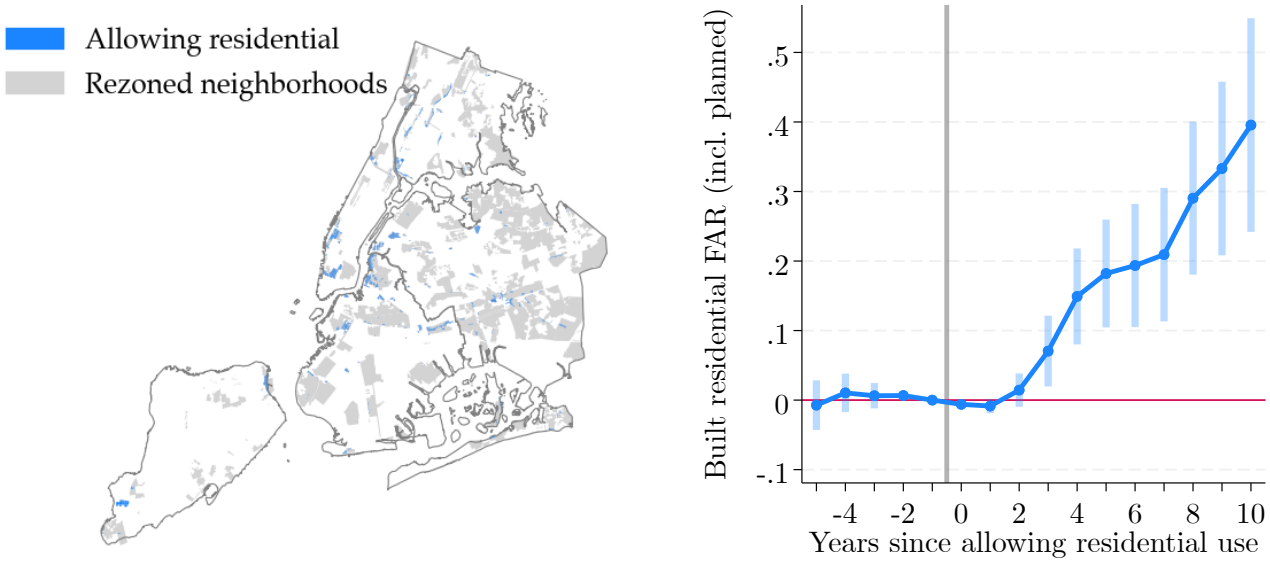
Amenity	**Correlation between share of satisfied residents and** B
Neighborhood cleanliness	0.49
Control of street noise	0.10
Household garbage pick-up	0.35
Recycling services	0.37
Snow removal	0.78
Rat control	0.39
Bike safety	0.28
Pedestrian safety	0.54
Street maintenance	0.53
Parking enforcement	0.74
Storm water drainage and sewer maintenance	0.58
Availability of healthcare services	0.53
Availability of cultural activities	0.61
Neighborhood parks	0.86
Fire protection services	0.76
Emergency medical services	0.74
Neighborhood public safety	0.67
Bus services	0.67
Subway services	0.73
Public services	0.63

Notes: This table compares calibrated amenities B with residents' satisfaction regarding 20 neighborhood characteristics, as measured in the 2017 Resident Survey conducted by the Citizens Budget Commission. I consider residents to be satisfied with an amenity if they rate it as "good" or "excellent" in their neighborhood. Survey results are available at the Community District level, a coarser geography than the Neighborhood Tabulation Areas used in the model. Therefore, I aggregate model-derived amenities B at the Community District level to compare them with survey results.

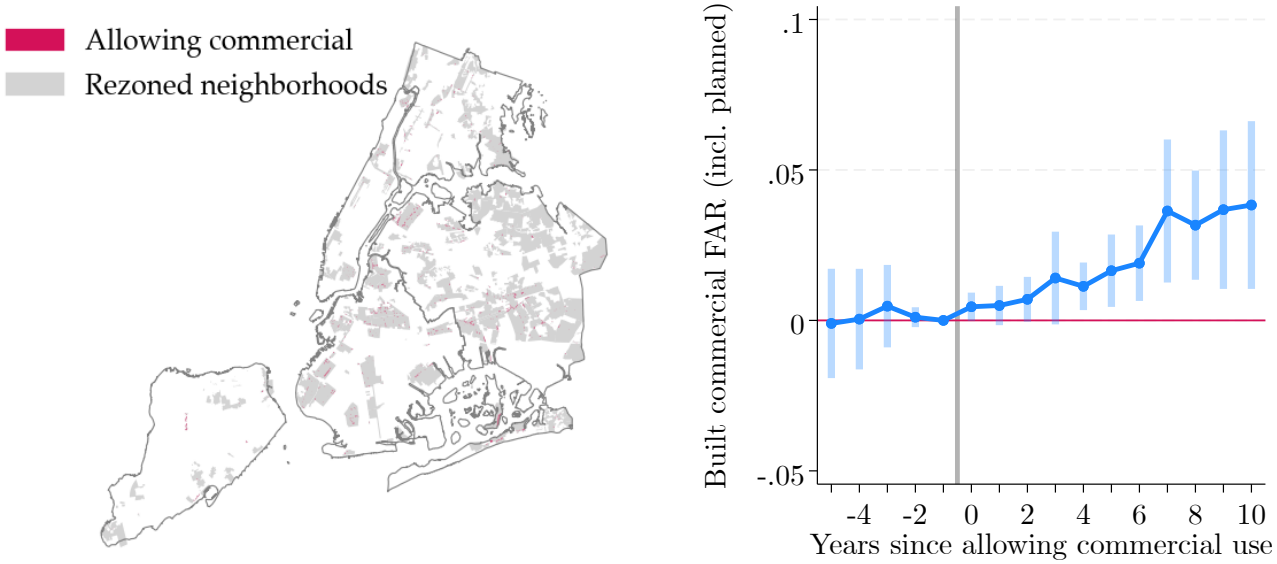
F Additional Results

Figure F.1: Effects of rezonings allowing previously restricted land uses

(a) Allowing residential uses

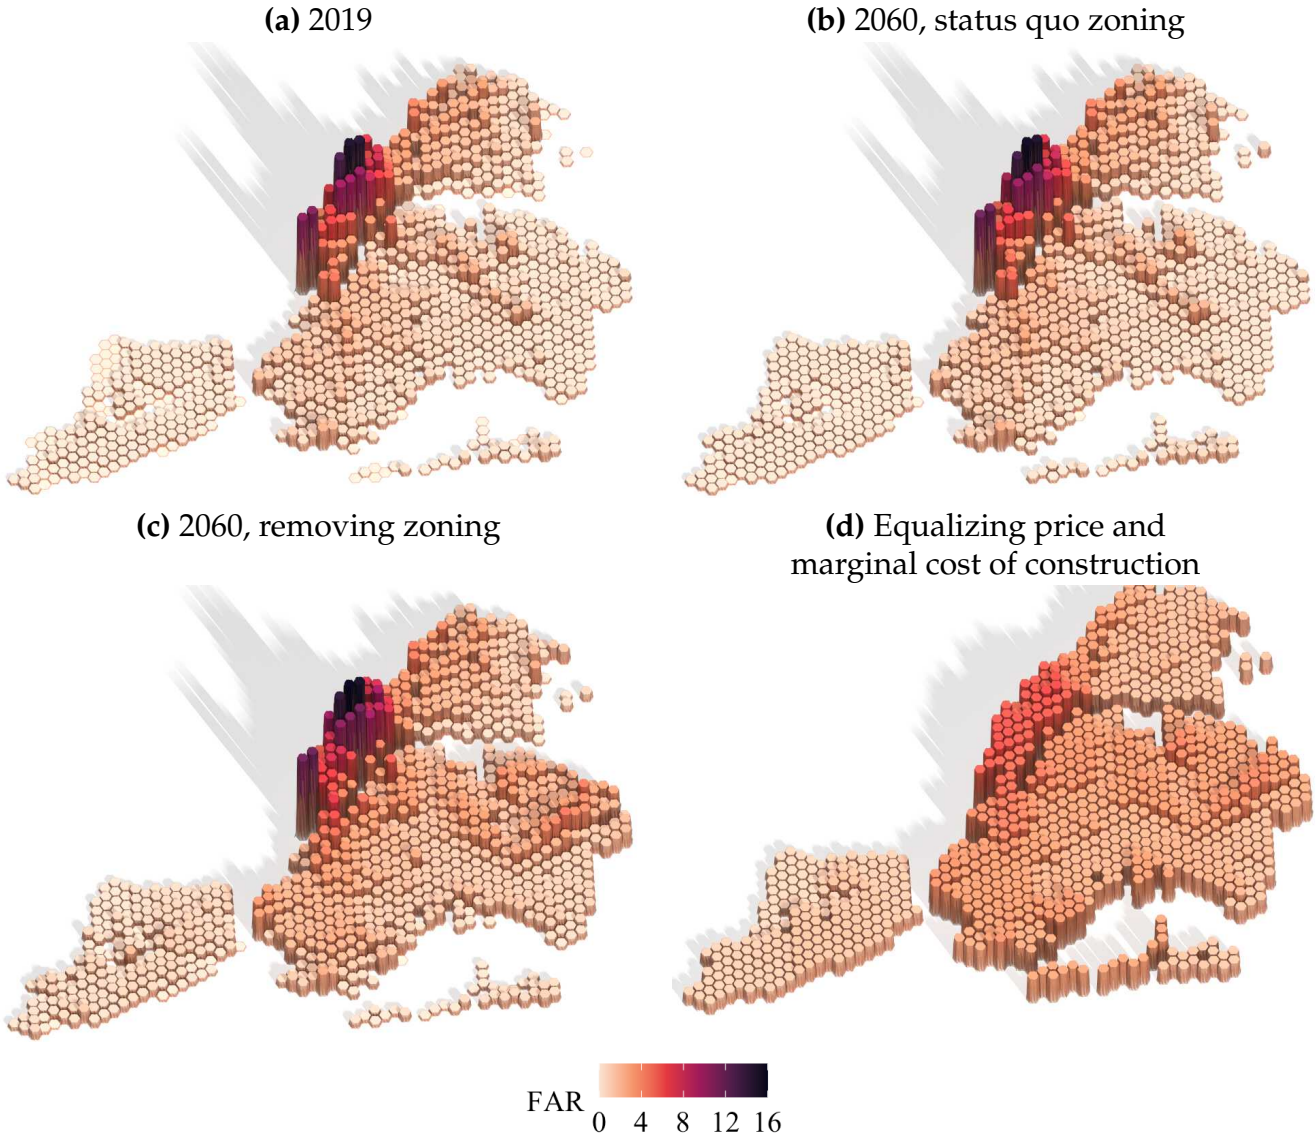


(b) Allowing commercial uses



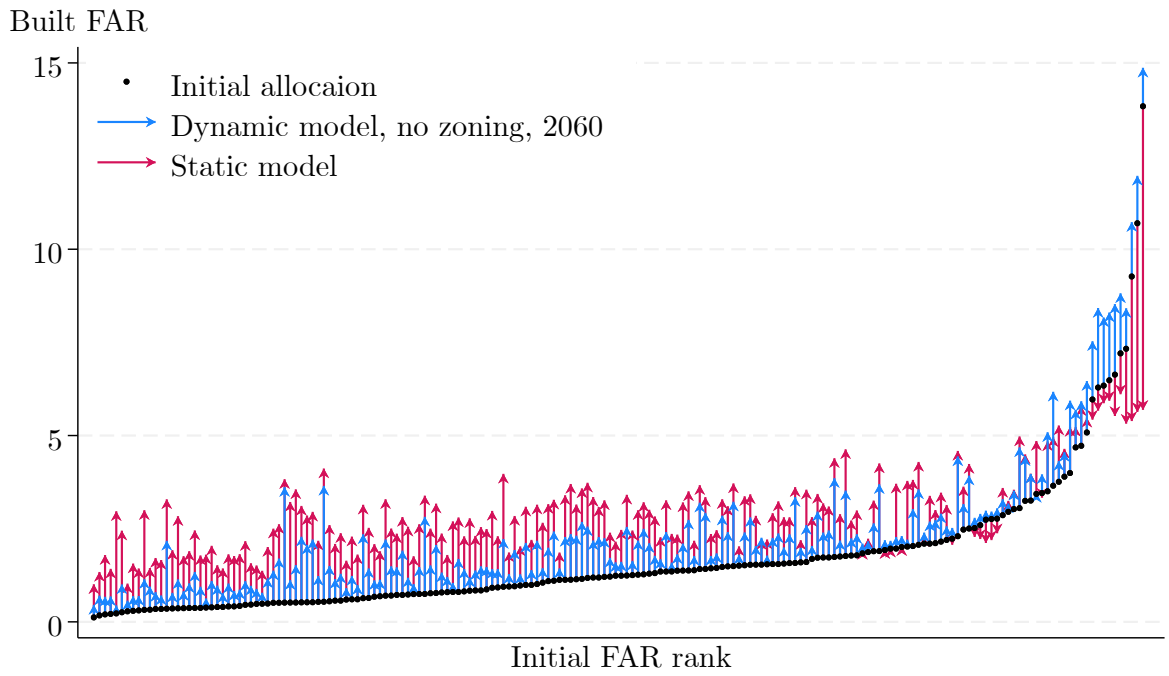
Notes: This figure shows the effects of rezonings allowing a previously restricted use. The estimation sample is composed of parcels within recently rezoned neighborhoods for which allowed uses changed. Panel (a) (resp., b) focuses on parcels for which rezoning allowed a previously disallowed residential (resp. commercial) use. The parcels in the estimation sample are mapped on the left panels, along with the rezonings that took place in NYC between 2004 and 2022. The right panels show the effect of rezoning on the built residential and commercial FAR (including the planned floorspace of buildings under construction). Treatment effects are computed using the procedure of De Chaisemartin and d'Haultfoeuille (2020).

Figure F.2: Counterfactual FAR profiles of NYC



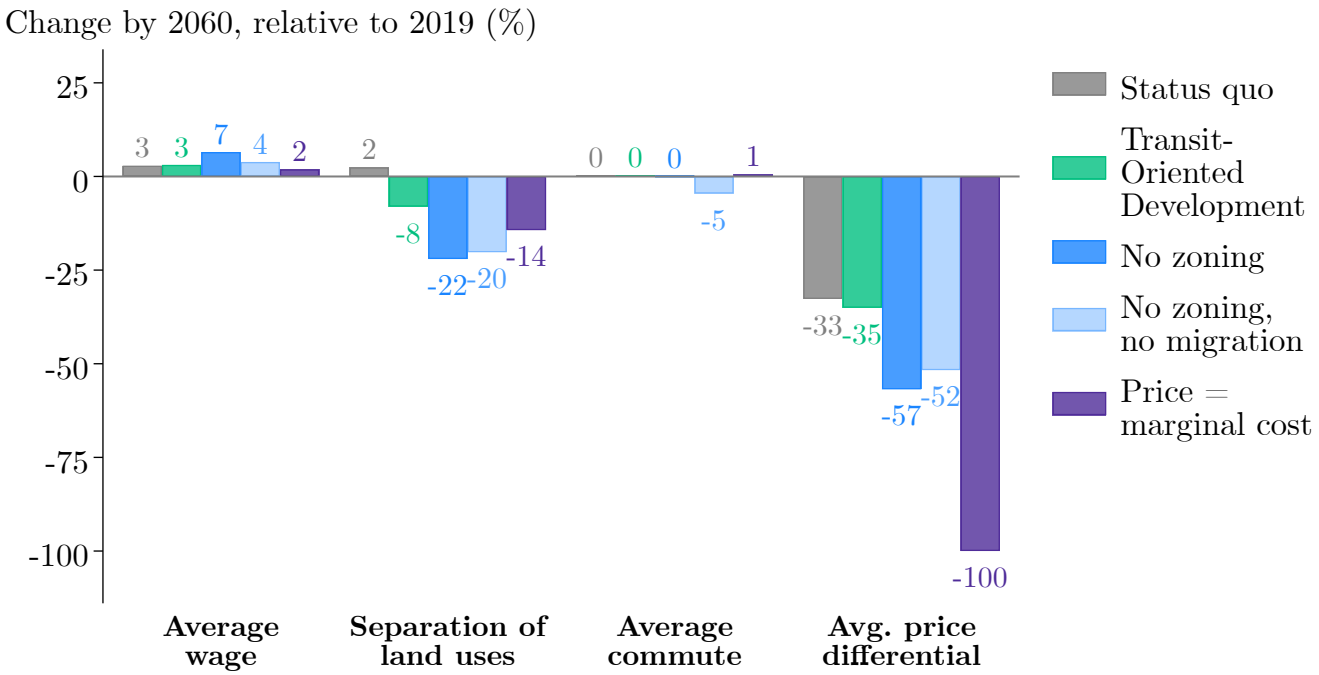
Notes: This figure shows the FAR profile of NYC for different counterfactual scenarios. Panel (a) shows FAR levels in 2019. Panel (b) (resp., c) shows expected FAR levels in 2060 if zoning were to remain at its 2019 level (resp., was removed). Panel (d) shows the distribution of FAR that equalizes the construction cost of the marginal sq. ft of floorspace with its price. In this figure, I aggregate data to hexagonal grid cells using Uber's H3 indexing system.

Figure F.3: Effects of removing zoning, computed using a dynamic vs. static approach



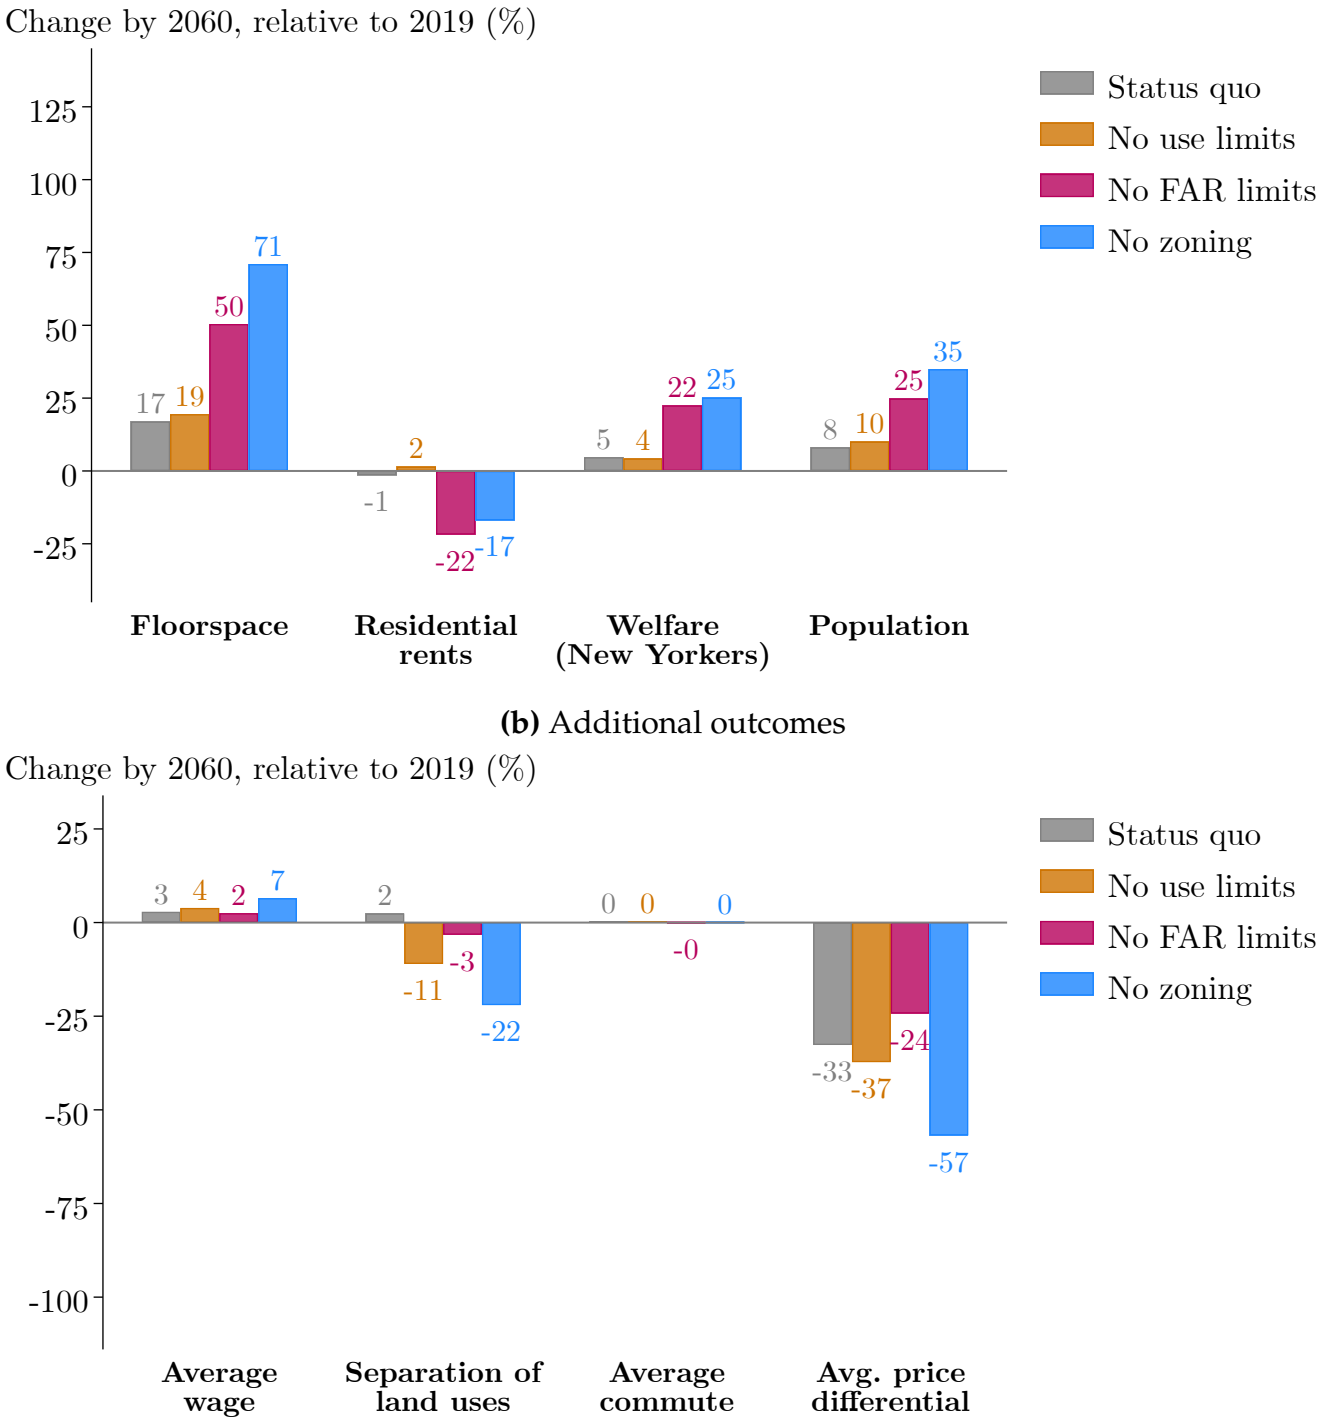
Notes: This figure shows the evolution of neighborhoods' FAR if zoning were removed, computed using two approaches. In blue, I show predictions of the dynamic model presented in this paper. In red, I show predictions of a static model, where the price of floorspace is equalized with its marginal construction cost in each neighborhood.

Figure F.4: Effects of relaxing zoning, additional outcomes



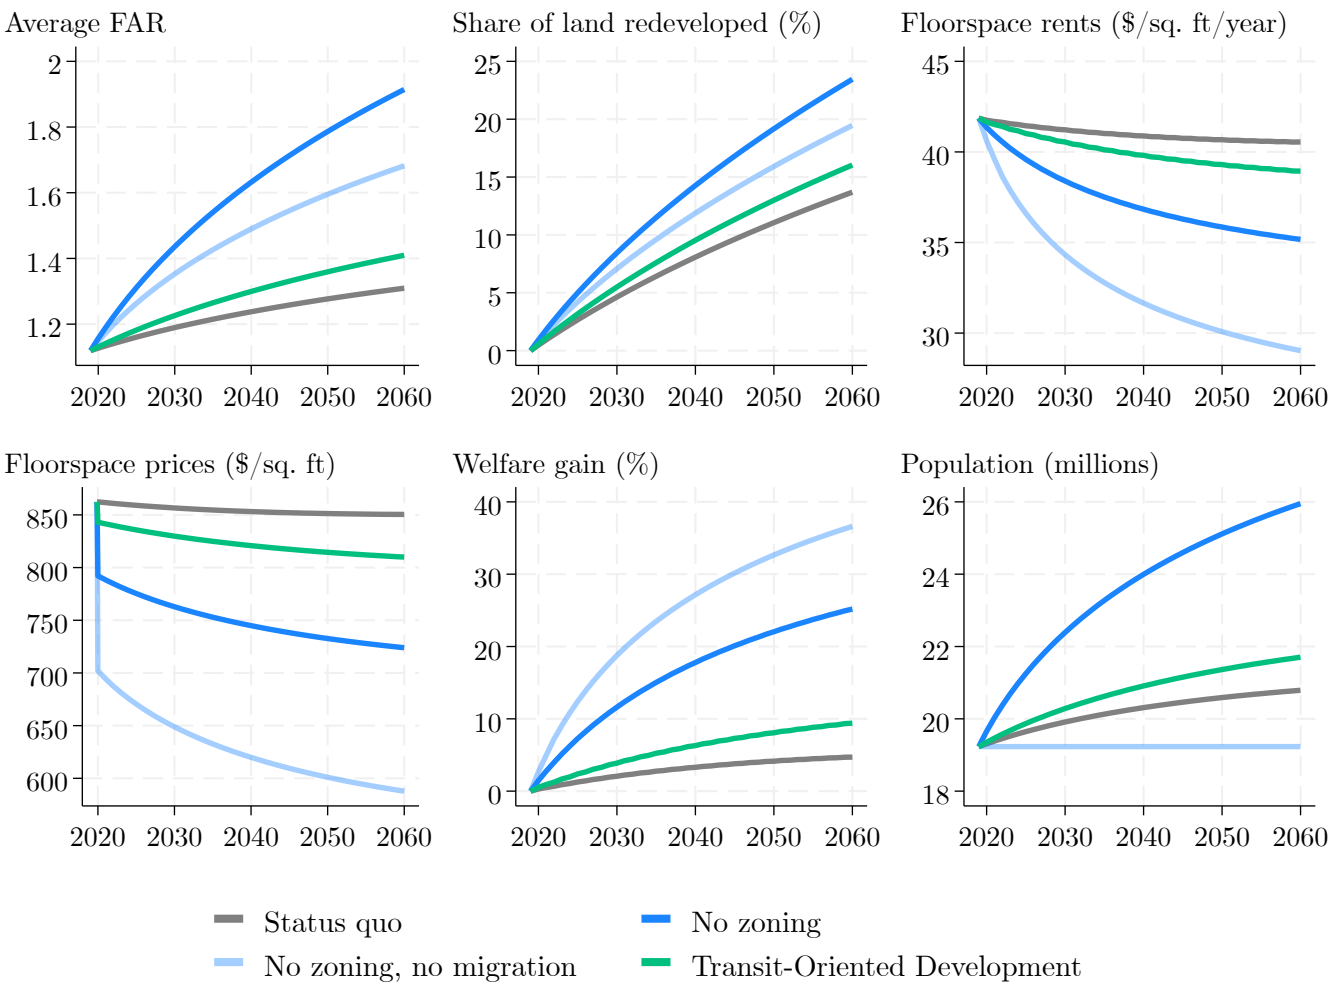
Notes: This figure shows changes in average wages, the separation of land uses, average commuting times, and the average residential–commercial price differential between 2019 and 2060 in the different counterfactuals described in Figure 7. The separation of land uses is measured through a dissimilarity index, as $S = 0.5 \times \sum_i |(H_{Ri}/H_R) - (H_{Fi}/H_F)|$, where H_{Ri} (resp., H_{Fi}) is the amount of residential (resp., commercial) floorspace in neighborhood i, and H_R (resp., H_F) is the total amount of residential (resp., commercial) floorspace in the city. Price differentials are measured as the average absolute difference between the price of residential floorspace and that of commercial floorspace, weighted by the initial amount of floorspace in each neighborhood.

Figure F.5: Effects of relaxing use vs. FAR limits



Notes: This figure shows changes in the outcomes studied in Figure 7 and Appendix Figure F.4 between 2019 and 2060 across four different counterfactuals. In the status quo counterfactual, zoning limits remain at their 2019 level. In the "no use limits" counterfactual, I allow residential and commercial uses on each parcel, keeping the maximum allowed FAR fixed. In the "no FAR limits" counterfactual, I remove FAR limits, but use limits remain at their 2019 level. In the "no zoning" counterfactual, I remove both use and FAR limits. In all counterfactuals, I keep zoning regulations at their 2019 levels for parcels protected as landmarks or part of historic districts or flood zones.

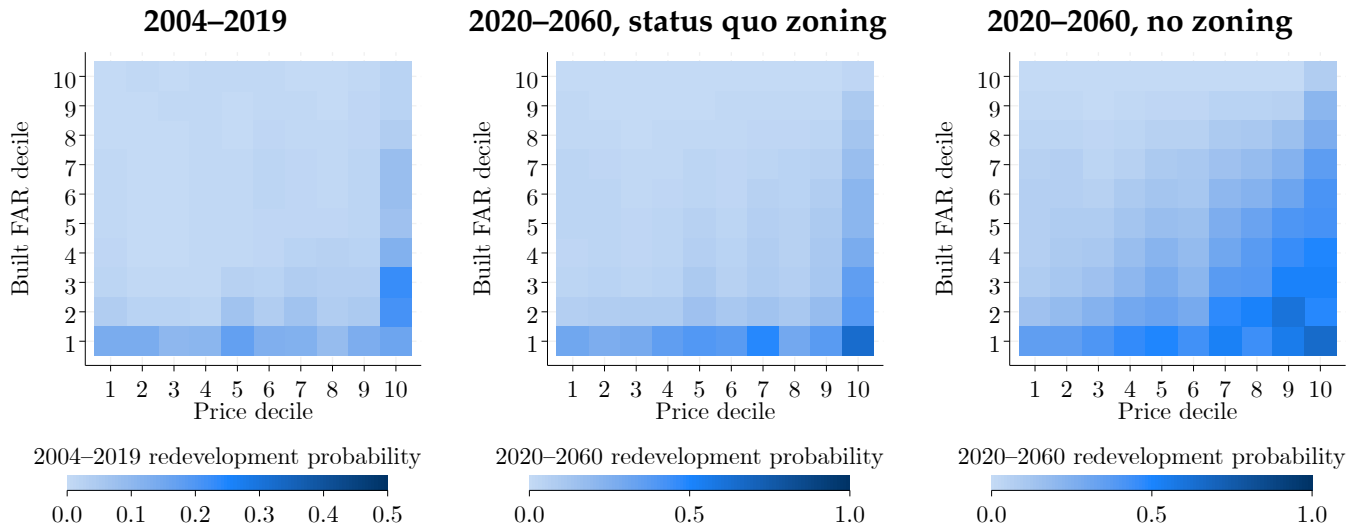
Figure F.6: Transition paths



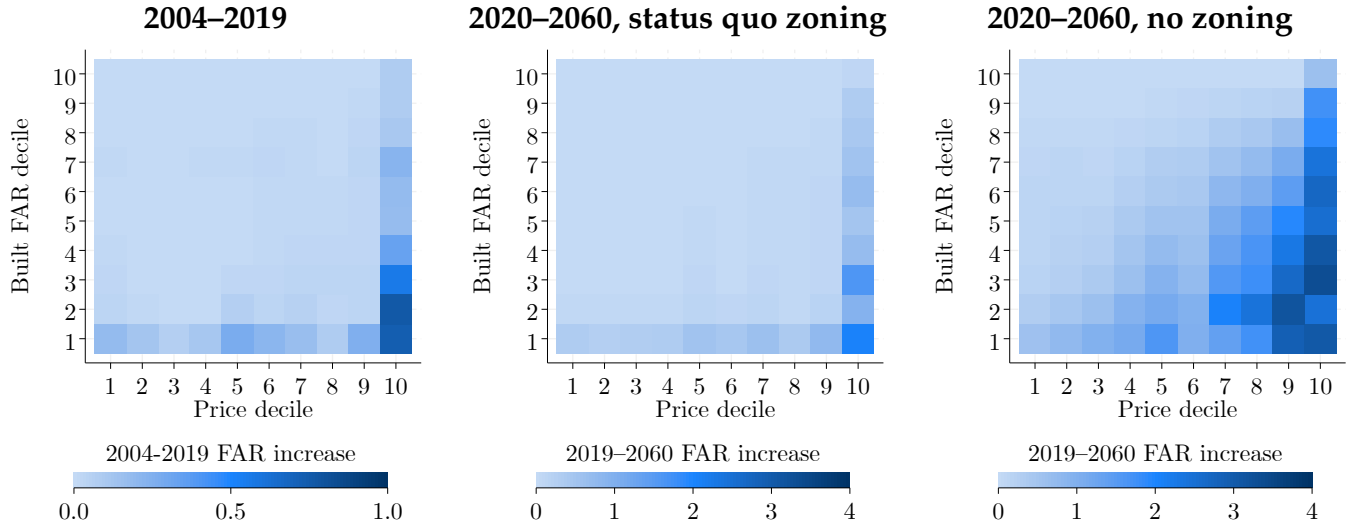
Notes: This figure shows how key variables evolve along the transition path. Floorspace prices (resp., rents) correspond to the average floorspace price (resp., rent) across neighborhoods, weighted by initial floorspace quantities.

Figure F.7: Redevelopment activity, split by price and density levels

(a) Redevelopment probabilities

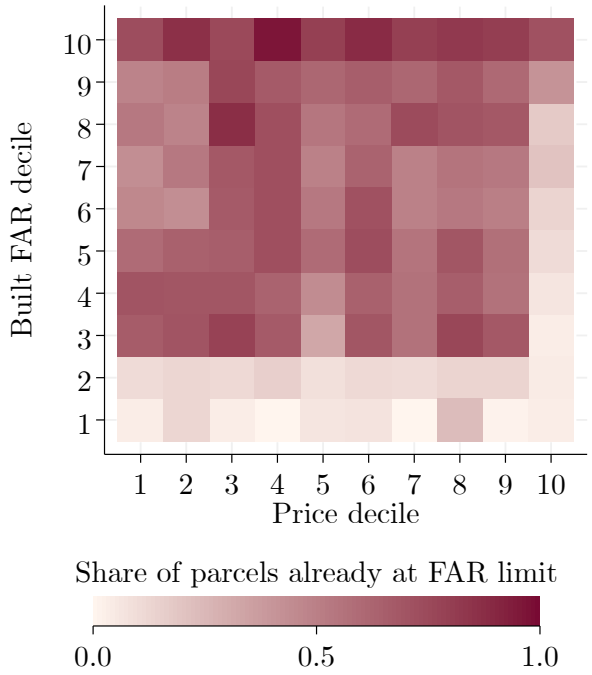


(b) FAR increases



Notes: This figure plots redevelopment activity for different deciles of floorspace prices and built FAR. Panel (a) shows redevelopment probabilities and panel (b) plots increases in built FAR. In each panel, I plot observed redevelopment activity over 2004–2019, predicted redevelopment activity over 2020–2060 if zoning stayed at its current level, and predicted redevelopment activity over 2020–2060 if zoning were removed.

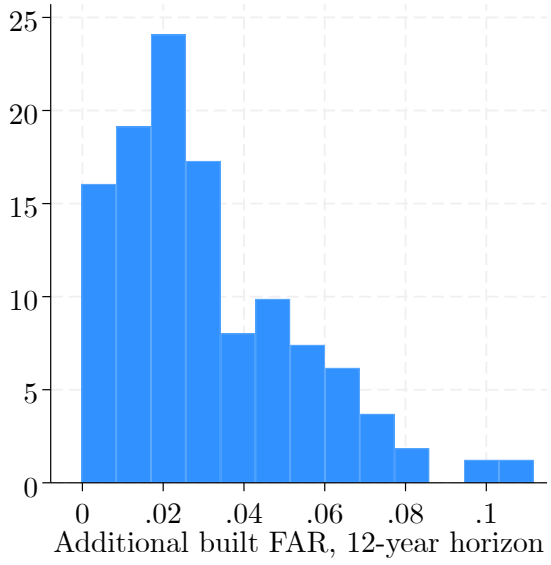
Figure F.8: Share of parcels with a built FAR at or above the zoning limit, split by price and density levels



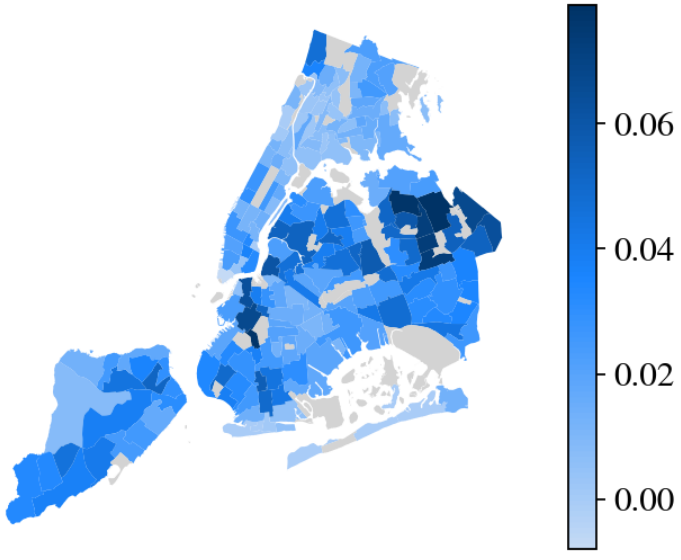
Notes: This figure plots the share of parcels with a built FAR over 80% of the maximum allowed FAR for different deciles of floorspace prices and built FAR.

Figure F.9: Effects of a uniform 1 FAR point upzoning on additional built FAR (10-year horizon)

(a) Heterogeneity in neighborhoods' growth

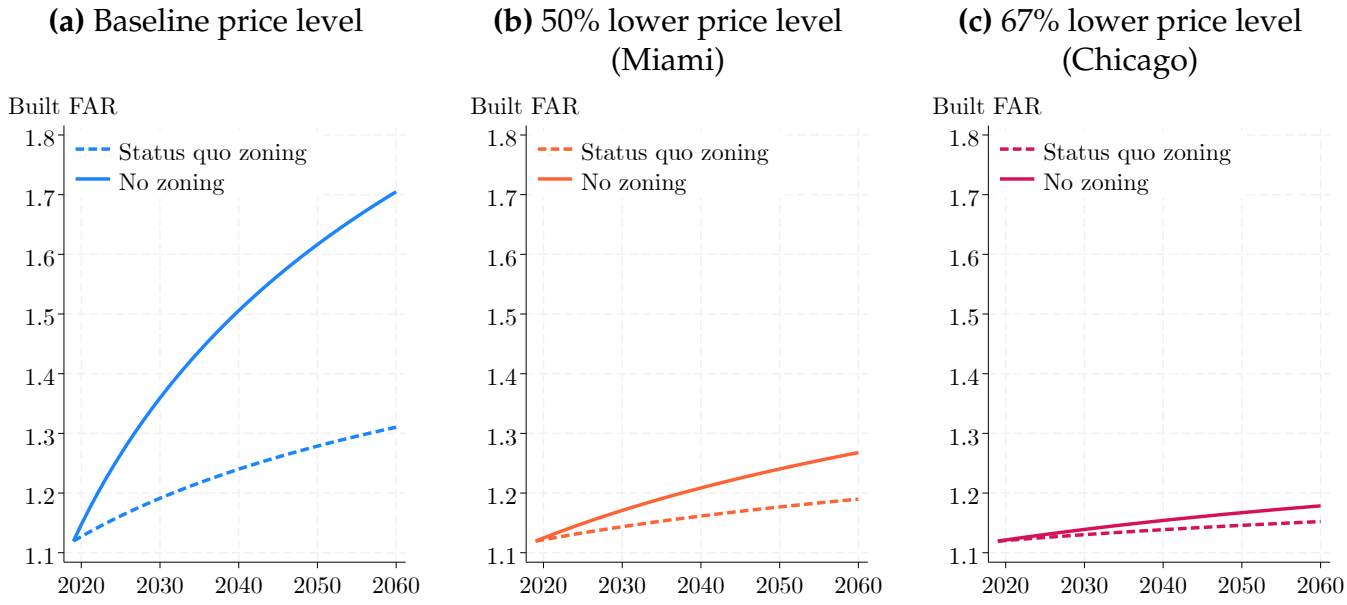


(b) Spatial distribution



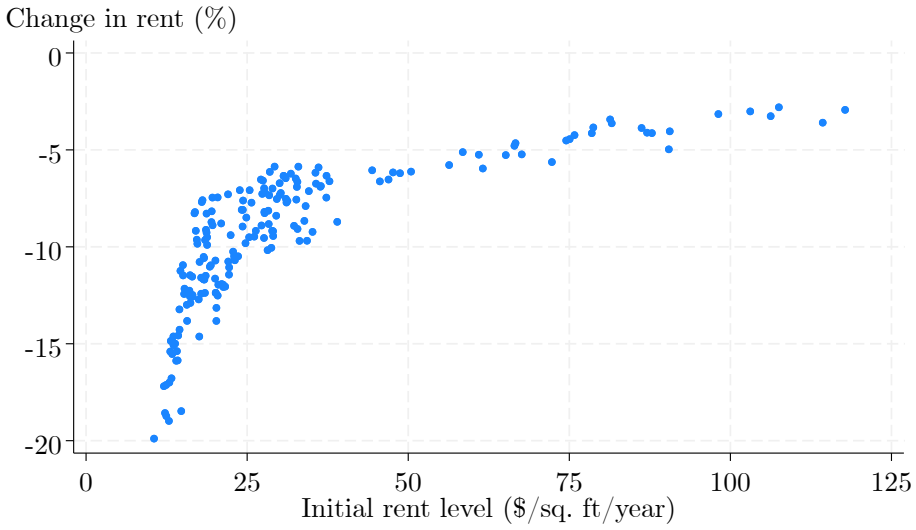
Notes: This figure shows the effect of uniformly upzoning NYC by one FAR point on the growth of different neighborhoods over a 10-year horizon.

Figure F.10: Effects of removing zoning on built FAR at different aggregate price levels



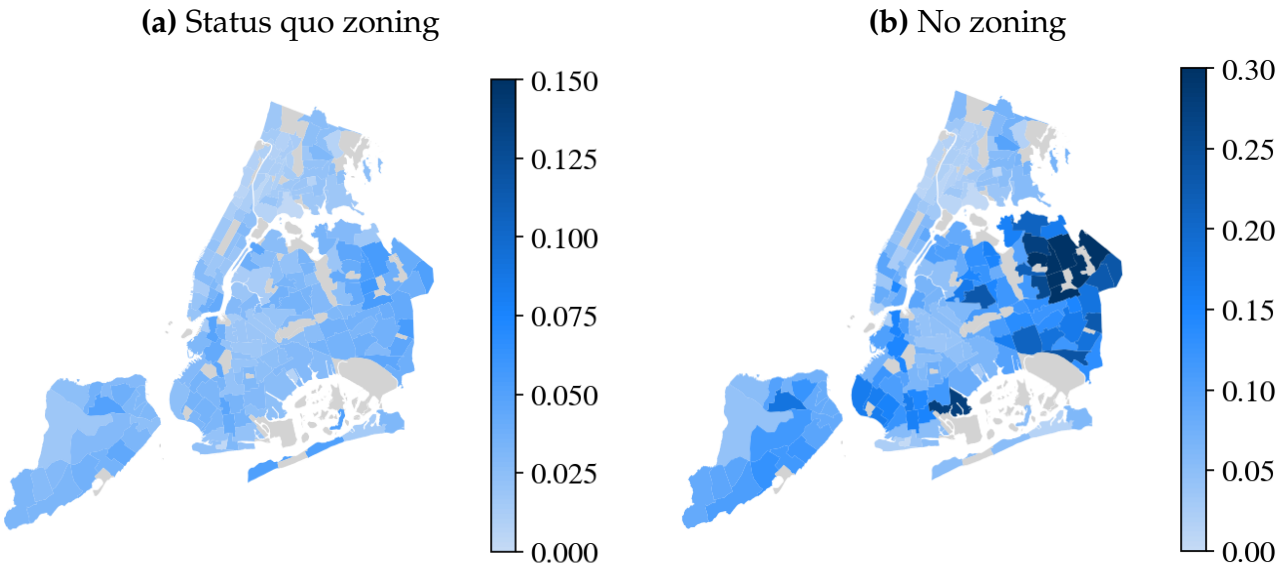
Notes: Panel (a) shows the evolution of the amount of floorspace in NYC under the baseline scenario, where productivity and amenity levels remain at their 2019 levels in the future. In Panels (b) and (c), I consider a scenario where the city is hit in 2020 with permanent productivity and amenity shocks that uniformly reduce floorspace prices (by 50% and 67%, respectively), placing them at levels comparable to those in Miami and Chicago, respectively.

Figure F.11: Effects of a 10% increase in the city's floorspace on rents



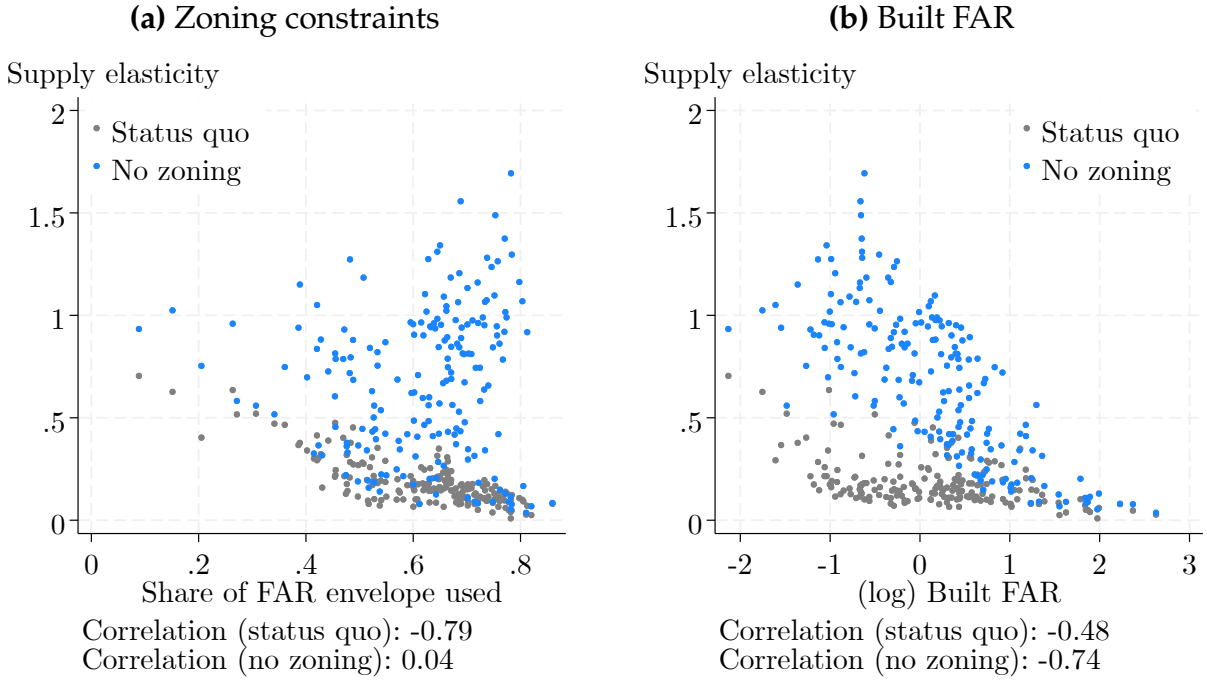
Notes: This figure plots the effects of uniformly increasing the city's floorspace by 10% on rents. Each dot corresponds to a neighborhood.

Figure F.12: Exposure to redevelopment



Notes: This figure maps, for each neighborhood, the share of residential floorspace in each neighborhood that will have been redeveloped by 2060 in the status quo zoning counterfactual (panel a) and the no zoning counterfactual (panel b).

Figure F.13: Correlates of neighborhood-level supply elasticities



Notes: This table shows the correlation between neighborhood-level supply elasticities (measured at a 40-year horizon) and two of their most important determinants: the share of the FAR envelope that is already used, measuring the extent to which zoning constraints are binding in each neighborhood (panel a), and preexisting density levels (panel b). Results are shown both under current zoning regulations and in the no-zoning counterfactual. The FAR envelope corresponds to the maximum amount of floorspace that could potentially be built according to the zoning code. When computing the share of the FAR envelope that is already used, I exclude from each parcel floorspace that exceeds the zoning code's FAR limit.