Affordable Housing and City Welfare *

Jack Favilukis[†]
UBC Sauder School of Business

Pierre Mabille[‡]
NYU Stern

Stijn Van Nieuwerburgh[§]
NYU Stern, NBER, and CEPR

February 15, 2018

Abstract

Housing affordability has become one of the main policy challenges for the major cities of the world. Two key policy levers are zoning and rent control policies. We build a new dynamic equilibrium asset pricing model to evaluate the implications of such policies for house prices, rents, production and income, residential construction, income and wealth inequality, as well as the spatial distribution of households within the city. We calibrate the model to New York City, incorporating current zoning and rent control systems. Our model suggests large welfare gains from relaxing zoning regulations in Manhattan, and more modest gains from reducing the size of the rent control program. The former policy is progressive and a Pareto improvement, while rent control reform is regressive in nature and hurts the current beneficiaries.

JEL classification: R10, R20, R30, R40, R51, G11, G12, H41, H70, J61
Keywords: Dynamic spatial equilibrium, house prices, affordable housing, rent control, zoning, gentrification, property taxes, real estate development

*First draft: February 15, 2018. We thank Hae-Kang Lee for excellent research assistance.

†Department of Finance, Sauder School of Business, University of British Columbia, 2053 Main Mall, Vancouver, BC, V6T 1Z2, Canada; jack.favilukis@sauder.ubc.ca; `https://sites.google.com/site/jackfavilukis`.

‡Department of Economics, Stern School of Business, New York University, 44 W. 4th Street, New York, NY 10012, USA; pmabille@stern.nyu.edu.

§Department of Finance, Stern School of Business, New York University, 44 W. 4th Street, New York, NY 10012, USA; svnieuwe@stern.nyu.edu; `http://www.stern.nyu.edu/~svnieuwe`.

1 Introduction

The increasing appeal of major urban centers has brought on an unprecedented housing affordability crisis. Ever more households are burdened by rents or mortgage payments that take up a large fraction of their paycheck as well as by long commutes. Local politicians are under pressure to improve affordability as well as to preserve socio-economic diversity in every part of the city. Economists have argued that zoning restrictions are making our most productive cities smaller than they should be, reducing macroeconomic growth, and are contributing to economic inequality and spatial segregation.

Policy makers have two main (categories of) tools to address housing affordability: rent control and zoning. Rent control creates affordable rental housing by government fiat. There are multiple policy levers: the rent discount to the market rent, the income threshold for qualifying, the size of the program, and its spatial distribution. Zoning governs land use and can be relaxed to increase the supply of housing, and all else equal, reduce its cost. Zoning changes are often associated with requirements on developers to set aside affordable housing units. Each policy affects the quantity and price of owned and rented housing and its spatial distribution. It affects incentives to work, wages, and commuting patterns, as well as property tax revenue. It also affects wealth inequality in the city and in each of its parts. While there is much work, both empirical and theoretical, on housing affordability, what is missing is a general equilibrium model that can quantify the impact of such policies on prices and quantities, on the spatial distribution of households, on income inequality within and across neighborhoods, and ultimately on individual and city-wide welfare. This paper sets up and solves a model that is able to address these questions.

We model a metropolitan area (city), which consists of two zones, the city center or central business district (zone 1) and the rest of the metropolitan area (zone 2). Working-age households who live in zone 2 commute to zone 1 for work. The model is an overlapping generations model with risk averse households that face labor income risk during their life-cycle and make dynamic decisions on consumption, savings, labor supply, tenure status (own or rent), and location. They face mortality risk in retirement. Since households cannot hedge labor income and longevity risk, markets are incomplete. This opens up the possibility that housing affordability policies provide insurance. We model a progressive tax system to capture the tax and transfer programs already in place. The model generates a rich cross-sectional distribution over age, labor income, tenure status, housing wealth, and financial wealth. This richness is paramount to understanding the distributional effects from affordability policies.

The city produces tradable goods and residential housing in each zone, subject to decreasing returns to scale. Firm profits are redistributed to local residents. Construction is subject to zoning policies which limit the maximum amount of housing that can be built, and thereby affect the housing supply elasticity. The city has a rent control system in place. Qualifying households can enter in a housing lottery that allocates housing units with steeply discounted rents. Rent control affects rents earned by landlords, which distorts the average price of rental buildings, and ultimately incentives to develop more residential space. While interest rates and goods prices are taken as given, wages, house prices, and market rents are determined in the city's equilibrium.

We calibrate the model to the New York metropolitan area. Our calibration matches key features of the data, including the relative size of Manhattan versus the rest of the metropolitan area, the New York metro income distribution, current zoning laws, the current size and scope of the rent control system, and current house prices and rents.

We use this model as a laboratory to explore various housing affordability policies. This allows us to highlight the channels through which rent control, an important feature of urban housing markets like New York and a policy pursued by other cities with renewed interest, alters the mechanisms of general equilibrium models with housing. Rent control provides an in-kind subsidy to households with a high marginal utility from housing services, creates tension between insiders and outsiders, and gives rise to a (spatial and sectoral) misallocation of labor and housing. The baseline model is calibrated to match the prevalence of rent control (RC) in each of the two zones in the New York metro, the observed discount of RC rents to market rents, and the typical size of a RC unit. The first set of policy experiments changes each of these policy levers one at the time, in order to study the welfare effects of reducing the generosity of the RC program.

Welfare results are nuanced. They depend on which policy lever is chosen, and sharply differ between homeowners and renters. The largest welfare gain comes from reducing the share of housing devoted to RC. Modestly spurring construction, the policy lowers market rents and improves housing affordability, as households spend a smaller fraction of income on rent. As we discuss below, this is consistent with the empirical literature about the effects of rent deregulation. Faced with lower rents, more households choose to live in Manhattan, lessening spatial segregation. Poorer households living in cheaper and smaller units coexist with richer households who benefit from the proximity to work.The biggest improvement in housing affordability shows up in the rest of the city. Income inequality in Manhattan increases, while metro-wide inequality falls. The policy hurts households who were in RC units prior to the policy change. It benefits market renters and owners who enjoy the improved housing affordability. More generally, this policy is regressive with welfare gains that increase in income and wealth. As a side effect, it generates spillovers to the market for owner-occupied housing where it spurs on home ownership.

A policy that reduces the maximum fraction of income that rent-controlled households are allowed to spend on rent, illustrates one of the trade-offs that determine social welfare. It brings into conflict insiders currently in RC and rent-burdened outsiders in market rentals. By effectively decreasing the size of RC units, this policy allows more households to live in cheaper housing, at the expense of previously rent-controlled agents who are forced to downsize.

In all experiments, welfare losses and gains are amplified with age, because of the life-cycle profile of income volatility and the shape of the tax and transfer system. The insurance value of RC increases with age. Policies scaling back RC programs generate a wedge between poor rent-controlled retirees and rich retired homeowners. More generally, the direct welfare losses of households benefiting from RC are traded-off against smaller gains for owners and market renters, who benefit from the indirect, general equilibrium effects improving quantities and prices. Because RC households represent a small share of the NY population, a utilitarian welfare function goes in the same direction as the general equilibrium effects if they are large enough. If not, then policy changes to the RC system end up hurting overall welfare. Policies such as lowering the rent subsidy per RC unit or tightening income qualification requirements allow to better target RC units to households in the bottom of the income distribution. However, they create losers who benefited from the system before, whose rent burden they increase. In that situation, lessening distortions alone is not enough to improve social welfare.

More generally, RC policies in place benefit lower-income and older households through two channels. First, they are in-kind subsidies to households with a high marginal utility for housing services. Second, they offer an insurance value, allowing households who face income risk to maintain a stable quantity of housing in all states of the world. Revealing households' preferences for decreasing their rent burden and for housing size rather than for choosing their preferred location, a policy that shifts all RC units in Manhattan to the rest of the metro area creates only a minor welfare gain.

The last policy experiment concerns zoning. The baseline model is calibrated to capture the relative population size of the two NYC zones. We make it easier to build in zone 1 by increasing the allowed maximum residential buildable area. We keep the current RC system in place. We find large welfare gains from this policy, the largest among all policies we consider. It is a Pareto improvement, benefiting RC renters, market renters, and owners about equally. It is also a progressive policy, benefiting the lower-income and lower-wealth households more. By increasing the equilibrium amount of housing in Manhattan, the policy lowers rents and prices, and increases affordability. It allows more households to live in Manhattan, thereby lowering their commute and increasing their leisure time.

Across all policies, aggregate effects are economically significant, albeit modest, and reflect the spatial and sectoral misallocation of labor. Because scaling back RC programs creates a boom in the construction sectors in the two zones, labor flows from the consumption sector to the latter. There, it is subject to starker decreasing returns and constrained by buildable areas (especially in zone 1), which induces an endogenous decrease in its aggregate productivity. As a result, aggregate output and profits decrease slightly. Eventually, those feed back into households' budget constraints and dampen the general equilibrium gains coming from improvements in prices and quantities.

Our work is at the intersection of the macro-finance and urban economics literatures. On the one hand, a large literature in finance solves partial-equilibrium models of portfolio choice between housing (extensive and intensive margin), financial assets, and mortgages.[1] More recent work in macro-finance has solved such models in general equilibrium, adding aggregate risk, endogenizing house prices and sometimes also interest rates.[2] Like the former literature, our model features a life-cycle and a rich portfolio choice problem. It aims to capture key quantitative features of observed wealth accumulation and home ownership over the life-cycle. Like the latter literature, house prices, rents, and wages are determined in equilibrium. Because we model one city, interest rates are naturally taken as given. Like the macro-finance literature, we aim to capture key features of house prices, income inequality, and wealth inequality.

On the other hand, a voluminous literature in urban economics studies the spatial location of households and firms in urban areas.[3] On the consumer side, this literature studies the trade-off between the commuting costs of workers, the housing prices they face, and the housing expenditures they make. These models tend to be static and households

[1] Early examples are Campbell and Cocco (2003), Cocco (2005) and Yao and Zhang (2004). A recent example is Berger, Guerrieri, Lorenzoni, and Vavra (2015). Davis and Van Nieuwerburgh (2015) provides a recent summary of this literature.

[2] E.g., Favilukis, Ludvigson, and Van Nieuwerburgh (2017) and Kaplan, Mitman, and Violante (2016). Davis and Van Nieuwerburgh (2015) provides a recent summary of this literature as well. One related study to ours is Imrohoroglu, Matoba, and Tuzel (2016) who study the effect of the 1978 passage of Proposition 13 which lowered property taxes in California. They find quantitatively meaningful effects on house prices, moving rates, and welfare. Our model adds a spatial dimension and aggregate risk but abstracts from housing transaction costs.

[3] Brueckner (1987) summarizes the Muth-Mills monocentric city model. Rappaport (2014) introduces leisure as a source of utility and argues that the monocentric model remains empirically relevant. Rosen (1979) and Roback (1982) introduce spatial equilibrium in a static setting.

tend to be risk neutral or have quasi-linear preferences.[4] The lack of risk, investment demand for housing, and wealth effects makes it hard to connect these spatial models to the finance literature.[5] Studying the welfare effects of housing affordability policies on the local economy requires a model with wealth effects. Our model studies spatial equilibrium within a city. Households are free to move across neighborhoods each period, rent or own, and choose how much housing to consume. We close the housing market in that local landlords who own more housing than they consume rent to other locals.

Because it is a heterogeneous-agent, incomplete-markets model, agents choices and equilibrium prices depend on the entire wealth distribution. Because of the spatial dimension, households' location is an additional state variable that needs to be kept track of. We use state-of-the-art methods to solve the model.[6] The resulting model is a new laboratory which can be used to explore many important questions like the impact of rent control and zoning laws on house prices, inequality, and housing affordability.

Finally, our model connects to a growing empirical literature that studies the effect of rent control and zoning policies on rents, house prices, and housing supply[7]. We replicate

[4]Van Nieuwerburgh and Weill (2010) solve a dynamic spatial equilibrium model with many cities and many household types. However, households have quasi-linear preferences. Recent work on spatial sorting across cities includes Behrens, Duranton, and Robert-Nicoud (2014) and Eeckhout, Pinheiro, and Schmidheiny (2014). Guerrieri, Hartley, and Hurst (2013) study house price dynamics in a city and focus on neighborhood consumption externalities, in part based on empirical evidence in Rossi-Hansberg, Sarte, and Owens (2010).

[5]Hizmo (2015) and Ortalo-Magné and Prat (2016) study a portfolio choice problem where households make a once-and-for-all location choice between cities. Conditional on the location choice, they are exposed to local labor income risk and make an optimal portfolio choice. They have constant absolute risk aversion preferences and consume at the end of life. The model features absentee landlords. The models are complementary to ours in that they solve a richer portfolio choice problem in closed-form, and have a location choice across cities. We solve a within-city location choice, but allow for preferences that admit wealth effects, and allow for consumption and location choice each period.

[6]We extend the approach of Favilukis et al. (2017), which itself extends Gomes and Michaelides (2008) and Krusell and Smith (1998) before that.

[7]Autor, Palmer, and Pathak (2014) and Autor, Palmer, and Pathak (2017) find that the elimination of the rent control mandate on prices in Cambridge increased the value of decontrolled units and neighboring properties in the following decade, by allowing constrained owners to raise rents and increasing the amenity value of those neighborhoods through housing market externalities. The price increase spurred new construction, increasing the rental stock. Symmetrically, Diamond, McQuade, and Qian (2017) show that the expansion of the rent control mandate in San Francisco led to a reduction in the supply of available housing, by decreasing owners' incentives to rent below market prices, paradoxically contributing to rising rents and the gentrification of the area. While beneficial to tenants in rent control, it resulted in an aggregate welfare loss. In contrast, Diamond and McQuade (forthcoming) find that the use of the Low Income Housing Tax Credit, a financial incentive for landlords to rent their properties below market prices to low-income tenants, leads to house price appreciation and decreasing segregation in low-income neighborhoods, thereby increasing welfare. As in our paper, the nature of the rent control policy and its distributional consequences are essential. Sieg and Yoon (2017) also stress the welfare gains from quantity- and price-based affordable housing policies in a model for NYC, where low-income households endogenously decide to search and wait for public housing or rent-stabilized apartments. However, they focus on tenants' individual choices without analyzing the aggregate effects of rent control in general equilibrium.

the main results of this literature, such as the modest local housing booms and the change in socioeconomic diversity that follow the removal of rent control, as well as its sharply different implications for the welfare of RC renters versus owners and renters.

The rest of the paper is organized as follows. Section 2 sets up the model. Section 3 describes the calibration to the New York metropolitan area. Section 4 discusses its main results and implications for quantities and prices, the distribution of households and affordability. Section 5 studies counterfactual policy experiments in which we vary the nature and strength of the rent control and zoning tools. Section 6 concludes.

2 Model

The model consists of two geographies, the "urban core" and the "periphery", whose union forms the "metropolitan area" or "city." The urban core, which we refer to as zone 1, is the central business district where all employment takes place. The periphery, or zone 2, contains the outer boroughs of the city as well as the suburban areas that belong to the metropolitan area. In the context of NYC, zone 1 is Manhattan, while zone 2 contains the other 24 counties that make up the metro area. While clearly an abstraction of the more complex production and commuting patterns in large cities, the monocentric city assumption captures the essence of commuting patterns (Rappaport, 2014). The two zones differ in size, as explained below. The city has a fixed population normalized to one.[8]

2.1 Households

Preferences The economy consists of overlapping generations of households. There is a continuum of households of a given age a. Each household maximizes a utility function u over consumption goods c, housing h, and labor supply n, and utility is allowed to depend on location ℓ and age a. The dependence on location allows us to capture the commuting time and amenity differences across locations. We use a Cobb-Douglas

[8]Future work could study interactions between affordability policies and net migration patterns in an open-city model. Such a model would need to take a stance on a reservation utility and moving cost that potentially differ by age, productivity, and wealth. The resulting proliferation of free parameters would pose a challenge to calibration. The empirical evidence for the New York metropolitan area, discussed in Appendix B.6, suggests that the zero net migration assumption we make fits the data well.

aggregator over c, h, and n:

$$U(c_t, h_t, n_t, \ell_t, a) = \frac{u(c_t, h_t, n_t, \ell_t, a)^{1-\gamma}}{1-\gamma}, \tag{1}$$
$$u(c_t, h_t, n_t, \ell_t) = \chi^{\ell,a}(c_t) c_t^{\alpha_c} h_t^{\alpha_h} \left(1 - n_t - \phi^{\ell}_{T,t}\right)^{1-\alpha_c-\alpha_h}.$$

The coefficient of relative risk aversion is γ. The time cost of commuting is ϕ^{ℓ}_{T}; the total (non-sleeping) hours in a period of time are normalized to 1. There also is a financial cost ϕ^{ℓ}_{F} associated with commuting from zone ℓ. We normalize the financial and time cost of commuting from within zone 1 to zero.

The term $\chi^{\ell,a}(c_t)$ is an age- and location-specific taste-shifter. Specifically, $\chi^{2,a}(c_t) = 1$, $\chi^{1,a}(c_t) = 1$ if $c_t < \underline{c}$, $\chi^{1,a}(c_t) = \chi^W > 1$ if $c_t \geq \underline{c}$ and $a < 65$, and $\chi^{1,a}(c_t) = \chi^R > 1$ if $c_t \geq \underline{c}$ and $a \geq 65$. The utility shifter creates a complementarity between living in zone 1 and high consumption levels. This modeling device captures that certain amenities such as high-end entertainment, restaurants, museums, or art galleries are concentrated in Manhattan. By assuming that the benefit from such amenities only accrue above a certain consumption threshold, we not only capture the luxury nature of those amenities but we also generate neighborhood consumption externalities as in Guerrieri et al. (2013). This force will help the model generate more concentration of high-income and high-wealth households in Manhattan. We allow the amenity value of zone 1 for rich people to differ for working-age (χ^W) and retirement-age households (χ^R). A special case of the model arise for $\chi^R = \chi^W = 1$ when the taste shifter is turned off. Another special case is $\underline{c} = 0$, which turns on the amenity value of Manhattan for all households, rich or poor.

There are two types of households in terms of subjective time discount factor. One group of households have a high degree of patience β^H while the rest have a low degree of patience β^L. This preference heterogeneity helps the model match observed patterns of home ownership and wealth accumulation over the life cycle. A special case of the model sets $\beta^H = \beta^L$.

Endowments A household's labor income depends on the number of hours worked n, the wage per hour worked W, a deterministic component $G(a)$ which captures the hump-shaped pattern in average labor income over the life-cycle, and an idiosyncratic labor productivity shock z.

There is an exogenous retirement age of 65. After retirement, households earn a pension which is the product of an aggregate component $\overline{\Psi}$ and an idiosyncratic component ψ^{a,z_t} which has cross-sectional mean of one. The idiosyncratic component reflects labor productivity during the working years. Labor income is taxed to finance the pension system. Denote the tax paid on an income level y_t by $T(y_t)$. The function $T(\cdot)$ captures a progressive tax and transfer scheme.

Households face mortality risk which depends on age, p^a. Although there is no intentional bequest motive, agents who die leave accidental bequests. We assume that the number of people who die with positive wealth leave a bequest to the same number of agents alive of ages 21 to 65. These agents are randomly chosen, with one restriction. Patient agents (β^H) only leave bequests to other patient agents and impatient agents (β^L) only leave bequests to other impatient agents.[9] Conditional on receiving a bequest, the size of the bequest $\widehat{b}_{t+1}$ is a draw from the relevant distribution (different for β^H and β^L types). Because housing wealth is part of the bequest and the house price depends on the aggregate state of the economy, the size of the bequest is stochastic. Agents know the distribution of bequests, conditional on β type. The model captures several features of real-world bequests: many households receive no bequest, bequests arrive later in life and at different points in time for different households, and there is substantial heterogeneity among bequest sizes for those who receive a bequest.[10]

Rent Control Every household in the model enters the rent control lottery every period. A household that wins the lottery has the option to move into a rent controlled unit in a zone assigned by the lottery, provided she qualifies.[11] A winning household retains the option to turn down the RC unit and rent or own the unit of its choice in the location of its choice on the free market. The RC unit rents at a fraction κ_1 of the free market rent. If the household accepts the lottery win, it must abide by two conditions: (i) its income must be below a cutoff expressed as a fraction κ_2 of area median income (AMI), and (ii) the rent paid on the RC unit must be below a fraction κ_3 of AMI. Both of these conditions are consistent with NYC rent regulation rules. The latter condition effectively restricts the maximum size of the RC unit. The probability of winning the lottery for each zone is set such that the demand of RC units in each zone is equal to the supply, discussed below. Households form beliefs about this probability. This belief must be consistent with rational expectations, and is updated as part of the numerical solution algorithm.

Since labor supply is endogenous in the model, a household may choose to reduce

[9] One interpretation is that attitudes towards saving are passed on from parents to children.

[10] As in Tabellini (1991), the young benefit from the unexpected house price appreciation of their owner parents through the type-specific bequest channel. Our bequest specification also captures that children have some idea about the kind of bequest they may expect to receive.

[11] There is a single lottery for all RC units. A certain lottery number range gives access to RC housing in zone 1, while a second range gives access to housing in zone 2. Households with lottery numbers outside these ranges lose the housing lottery.

hours worked in order to qualify for a RC unit. This has adverse implications for city-wide production. Also, a household may choose to consume more or less housing, conditional on being in a RC unit, than it would on the free market. This may lead to misallocation of housing in the cross-section of households.

Location and Tenure Choice Denote by $p^{RC,\ell}$ the probability of winning the RC lottery and being offered a unit in zone ℓ. The household chooses whether to accept the RC option with value $V_{RC,\ell}$, or to turn it down and go to the private housing market with value V_{free}. The value function V is:

$$V = p^{RC,1} \max\left\{V_{RC,1}, V_{free}\right\} + p^{RC,2} \max\left\{V_{RC,2}, V_{free}\right\} + \left(1 - p^{RC,1} - p^{RC,2}\right) V_{free}.$$

A household that loses the lottery or wins it but turns it down, freely chooses in which location ℓ to live and whether to be an owner (O) or a renter (R).

$$V_{free} = \max\left\{V_{O,1}, V_{R,1}, V_{O,2}, V_{R,2}\right\}.$$

The Bellman equations for $V_{RC,\ell}$, $V_{R,\ell}$ and $V_{O,\ell}$ are defined below.

Let S_t be the aggregate state of the world, which includes the wage W_t, as well as the housing price P_t^ℓ, the market rent R_t^ℓ and previous housing stock H_{t-1}^ℓ for each location ℓ. The household's individual state variables are its net worth at the start of the period x_t, its idiosyncratic productivity level z_t, and its age a. We suppress the dependence on β-type in the problem formulation below, but note here that there is one set of Bellman equations for each β-type.

Market Renter Problem A renter household on the free rental market in location ℓ chooses non-durable consumption c_t, housing consumption h_t, and working hours n_t to solve:

$$\begin{aligned}
&V_{R,\ell}(x_t, z_t, a, S_t) = \max_{c_t, h_t, n_t} U(c_t, h_t, n_t, \ell_t) + (1 - p^a)\beta\mathbb{E}_t[V(x_{t+1}, z_{t+1}, a+1, S_{t+1})] \text{ s.t.}\\
&c_t + R_t^\ell h_t + Q_t b_{t+1} + \phi_{F,t}^\ell = y_t - T(y_t) + \overline{\Psi}_t \psi^z + x_t + \Pi_t\\
&y_t = W_t n_t G^a z_t\\
&x_{t+1} = b_{t+1} + \widehat{b}_{t+1} \geq 0\\
&n_t \geq 0
\end{aligned}$$

The renter's savings in the risk-free bond, $Q_t b_{t+1}$, are obtained from the budget constraint. Pre-tax income y_t is the product of wages W per efficiency unit of labor times the number of efficiency units of labor. The latter is the product of the number of hours $n \geq 0$ times the productivity per hour $G^a z$. G^a is a deterministic age component while z is a stochastic, idiosyncratic, persistent productivity component. Households receive profits from local firms Π_t, rebated lump-sum. Next period's financial wealth x_{t+1} consists of savings b_{t+1} plus any accidental bequests $\widehat{b}_{t+1}$. Housing is divisible and there are no moving costs.[12]

RC Renter Problem A renter household in the RC system in location ℓ chooses non-durable consumption c_t, housing consumption h_t, and working hours n_t to solve:

$$
\begin{aligned}
&V_{RC,\ell}(x_t, z_t, a, S_t) = \max_{c_t,h_t,n_t} U(c_t, h_t, n_t, \ell_t) + (1 - p^a)\beta \mathbb{E}_t[V(x_{t+1}, z_{t+1}, a + 1, S_{t+1})] \text{ s.t.} \\
&c_t + \kappa_1 R_t^\ell h_t + Q_t b_{t+1} + \phi_{F,t}^\ell = y_t - T(y_t) + \overline{\Psi}_t \psi^z + x_t + \Pi_t \\
&x_{t+1} = b_{t+1} + \widehat{b}_{t+1} \geq 0 \\
&n_t \geq 0 \\
&y_t \leq \kappa_2 \overline{Y}_t \\
&\kappa_1 R_t^\ell h_t \leq \kappa_3 \overline{Y}_t
\end{aligned}
$$

The rent (per square foot) of a RC unit is a fraction κ_1 of the market rent. The last two inequalities reflect the qualification requirements for RC spelled out above, where $\overline{Y}_t = Median[W_t G^a n_t^i z_t^i]$ is the metro area median income (AMI).[13]

Owner's Problem An owner in location ℓ chooses non-durable consumption c_t, housing consumption h_t, working hours n_t, and investment property $\widehat{h}_t$ to solve:

$$
\begin{aligned}
&V_{O,\ell}(x_t, z_t, a, S_t) = \max_{c_t,h_t,\widehat{h}_t,n_t} U(c_t, h_t, n_t, \ell_t) + (1 - p^a)\beta E_t[V(x_{t+1}, z_{t+1}, a + 1, S_{t+1})] \text{ s.t.} \\
&c_t + P_t^\ell h_t + Q_t b_{t+1} + \kappa_4^\ell P_t^\ell \widehat{h}_t + \phi_{F,t}^\ell = y_t - T(y_t) + \overline{\Psi}_t \psi^z + x_t + \Pi_t + \kappa_4^\ell R_t^\ell \widehat{h}_t \\
&x_{t+1} = b_{t+1} + \widehat{b}_{t+1} + P_{t+1}^\ell (h_t + \kappa_4^\ell \widehat{h}_t)(1 - \delta - \tau^P) \\
&-Q_t b_{t+1} \leq \theta_{res} P_t^\ell h_t + \theta_{inv} \kappa_4^\ell P_t^\ell \widehat{h}_t \\
&n_t \geq 0 \\
&\widehat{h}_t \geq 0 \\
&\kappa_4^\ell = 1 - \eta^\ell + \eta^\ell \kappa_1
\end{aligned}
$$

Local home owners are the landlords to the local renters. For simplicity, we assume

[12] Moving costs can be added but only at the expense of additional numerical complexity. Moreover, urban households are fairly mobile over periods of four years, which is the length of a period in our calibration.

[13] In the implementation, we assume that the income and size qualification cutoffs for RC are constants. We then compute what fractions κ_2 and κ_3 of AMI they represent. This allows us to sidestep the issue that the AMI may change with RC policies.

that renters cannot buy investment property and that owners can only buy investment property in the location of their primary residence. Owners earn rental income on their investment units.

Landlords in the model are required to buy η^ℓ square feet of rent controlled property for every $1-\eta^\ell$ square feet of market rental property. This captures the institutional reality of affordable housing programs in NYC (and elsewhere).[14] The effective rent earned per square foot of investment property is $\kappa_4^\ell R_t^\ell$. Since the regulated rent is a multiple κ_4 of the market rent, the regulated price must be the same multiple of the market price. Because prices and rents scale by the same constant, the return on investing in regulated units is identical to the return on investing in market units. As a result, landlords are unaffected by rent regulation. However, the lower average price for rental property ($\kappa_4 < 1$) has important effects on rental housing supply/development as discussed below.

The physical rate of depreciation for all housing units is δ. The term $Ph\delta$ is a financial costs, i.e., a maintenance cost.[15] As shown in equation (7) below, the physical depreciation $Ph\delta$ can be offset by residential investment undertaken by the construction sector.[16]

Property taxes on the housing owned in period t are paid in year $t+1$; the tax rate is τ^P. Property tax revenue is used for local government spending which does not confer utility to the households.[17]

Housing serves as a collateral asset for debt. For simplicity, mortgages are negative short-term safe assets. Households can borrow a fraction θ_{res} of the market value of their primary residence and a potentially different fraction θ_{inv} against investment property. The empirically relevant case is $\theta_{res} \geq \theta_{inv}$.

In the appendix we show that for renters, the choices of h_t and n_t are analytic functions of c_t, therefore the renter's problem can be rewritten with just two choices: consumption

[14]Examples of incentives provided for the development of affordable housing in NYC are (i) the 80/20 new construction housing program, a state program that gives low-cost financing to developers who set aside at least 20% of the units in a property for lower-income families; (ii) the 421a program, which gives tax breaks (up to 25 years) for the development of under-utilized or vacant sites often conditional on providing at least 20% affordable units (i.e., used in conjunction with the 80-20 program), (iii) the Federal Low Income Housing Tax Credits (LITCH) program, which gives tax credits to developers directly linked to the number of low-income households served, and (iv) Mandatory Inclusionary Housing, a New York City program that lets developers build bigger buildings and gives them tax breaks if they reserve some of the units for (permanently) affordable housing.

[15]It is easy to allow for an additional maintenance cost for rental housing ($\delta_{inv} > \delta$). This would make renting less attractive and increases the home ownership rate. We do not need a depreciation wedge to generate the observed home ownership rate.

[16]This treatment of depreciation avoids having to keep track of the aggregate owner-occupied fraction of housing as an additional state variable.

[17]Alternatively, one could solve a model where property tax revenue finances a local public good that enters the local residents' utility function. That requires taking a stance on the intra-temporal rate of substitution between private and public consumption.

c_t and location ℓ. For owners, the choices of h_t and n_t are analytic functions of c_t and $\widehat{h}_t$, therefore the owner's problem can be rewritten with just three choices: consumption c_t, investment property size $\widehat{h}_t$, and location ℓ.

2.2 Firms

Goods Producers There are a large number n_f of identical, competitive firms located in the urban core (zone 1), all of which produce the numéraire consumption good.[18] This good is traded nationally; its price is unaffected by events in the city and normalized to 1. The firms have decreasing returns to scale and choose labor inputs to maximize profit each period:

$$\Pi_{c,t} = \max_{N_{c,t}} N_{c,t}^{\rho_c} - N_{c,t} W_t \tag{2}$$

We assume that these firms are owned by local equity owners. For simplicity, we redistribute all profits in lump-sum fashion to local residents.

Developers and Affordable Housing Mandate In each location ℓ there is a large number n_f of identical, competitive construction firms (developers) which produce new housing units and sell them locally. For simplicity, we assume that all developers are headquartered in the urban core, regardless of where their construction activity takes place. As we assumed for the consumption good firms, construction firms are owned by local equity owners who receive their profits in lump-sum fashion.[19]

The cost of the affordable housing mandate is born by developers. The affordable housing regulation stipulates that for every square foot of market rental units built in zone ℓ, η^ℓ square feet of RC units must be built. Developers receive an average price per foot for rental property of $\kappa_4 P_t$, while they receive a price per foot of P_t for owner-occupied units. Given a home ownership rate in zone ℓ of ho_t^ℓ, developers receive an average price per foot of:

$$\overline{P}_t^\ell = \left(ho_t^\ell + (1 - ho_t^\ell)\kappa_4^\ell \right) P_t^\ell. \tag{3}$$

In sum, the cost of construction of owner-occupied and rental property in a given location is the same. After completion of construction but prior to sale, some of the newly constructed housing units are designated as rental units and the remainder as owner-

[18]We assume that the number of firms is proportional to the number of households in the city when solving the model. With this assumption, our numerical solution is invariant to the number of households. Due to decreasing return to scale, the numerical solution would depend on the number of households otherwise.

[19]Indeed, many housing developers in NYC are local owner-operator-developers.

ship units. The rental designation triggers affordable housing regulation, as reflected in a lower rent and price. Of course, developers would like to sell ownership units rather than rental units, but the home ownership rate is determined in equilibrium. They are price takers in the market for space and face an average price of $\overline{P}_t^\ell$.

A special case of the model is the case without rent control: $\kappa_4^\ell = 1$ either because $\eta^\ell = 0$ or $\kappa_1 = 1$. In that case, $\overline{P}_t^\ell = P_t$. Without rent control, the higher sale price for housing increases incentives to develop more housing.

Zoning Given the existing housing stock in location ℓ, H_{t-1}^ℓ, construction firms have decreasing returns to scale and choose labor to maximize profit each period:

$$\Pi_{h,t}^\ell = \max_{N_{\ell,t}} \left(1 - \frac{H_{t-1}^\ell}{\overline{H^\ell}}\right) \overline{P}_t^\ell N_{\ell,t}^{\rho_h} - N_{\ell,t} W_t \tag{4}$$

The production function of housing has two nonlinearities. First, as for consumption good firms, there are decreasing returns to scale because $\rho_h < 1$.

Second, construction is limited by zoning laws. The maximal amount of square footage zoned for residential use in zone ℓ is given by $\overline{H^\ell}$.[20] This term captures the idea that, the more housing is already built in a zone, the more expensive it is to build additional housing. For example, additional construction may have to take the form of taller structures, buildings on less suitable terrain, or irregular infill lots. Therefore, producing twice as much housing requires more than twice as much labor. Laxer zoning policy, modeled as a larger $\overline{H^\ell}$, makes development cheaper, and all else equal, will expand the supply of housing.

When $\overline{H^\ell}$ is sufficiently high, the model's solution becomes independent of $\overline{H^\ell}$, and the supply of housing is governed solely by ρ_h. When $\overline{H^\ell}$ is sufficiently low, the housing supply depends on both $\overline{H^\ell}$ and ρ_h.[21]

Per capita profits from tradeable and construction sectors are:

$$\Pi_t = \Pi_{c,t} + \Pi_{h,t}^1 + \Pi_{h,t}^2$$

and are reflected in households' budget constraints above.

[20] We interpret $\overline{H^\ell}$ as the total land area zoned for residential use multiplied by the maximum permitted number of floors that could be built on this land, the floor area ratio (FAR).

[21] In this sense, the model captures that construction firms must pay more for land when land is scarce or difficult to build on due to regulatory constraints. This scarcity is reflected in equilibrium house prices.

2.3 Equilibrium

Given parameters, a competitive equilibrium is a price vector $(W_t, P_t^\ell, R_t^\ell)$ and an allocation, namely aggregate residential demand by market renters $H_t^{R,\ell}$, RC renters $H_t^{RC,\ell}$, and owners $H_t^{O,\ell}$, aggregate investment demand by owners $\widehat{H}_t^\ell$, aggregate housing supply, aggregate labor demand by goods and housing producing firms $(N_{c,t}, N_{\ell,t})$, and aggregate labor supply N_t such that households and firms optimize and markets clear.

The following conditions characterize the equilibrium. First, given wages and average prices given by (3), firms optimize their labor demand, resulting in the first-order conditions:

$$N_{c,t} = \left(\frac{\rho_c}{W_t}\right)^{\frac{1}{1-\rho_c}} \quad \text{and} \quad N_{\ell,t} = \left(\frac{\left(1 - \frac{H_{t-1}^\ell}{\overline{H^\ell}}\right)\overline{P}_t^\ell \rho_h}{W_t}\right)^{\frac{1}{1-\rho_h}}. \tag{5}$$

Second, labor markets clear:

$$n_f \left(N_{c,t} + \sum_\ell N_{\ell,t}\right) = N_t. \tag{6}$$

Third, the housing market clears in each location ℓ:

$$(1-\delta)H_{t-1}^\ell + n_f \left(1 - \frac{H_{t-1}^\ell}{\overline{H^\ell}}\right) N_{\ell,t}^{\rho_h} = H_t^{O,\ell} + \widehat{H}_t^\ell. \tag{7}$$

The left-hand-side is the supply of housing which consists of the non-depreciated housing stock and new residential construction. The right-hand-side is the demand for those housing units by owner-occupiers and landlords.

Fourth, the supply of rental units in each location ℓ must equal the demand from market tenants and RC tenants:

$$\widehat{H}_t^\ell = H_t^{R,\ell} + H_t^{RC,\ell} \tag{8}$$

Fifth, average pension payments equal to average labor income taxes collected:

$$\overline{\Psi} N_{ret} = \int_i T(y_t^i) di, \tag{9}$$

where N_{ret} is the total number of retirees, which is a constant. The right-hand side is also constant in the steady state of the model.

Sixth, the aggregate state S_t evolves according to rational expectations.

Seventh, the value of all bequests received is equal to the wealth of all agents who die.

2.4 Welfare Effects of Affordability Policies

We compute the welfare effect of an affordability policy using the following procedure. Denote agent i's value function under benchmark policy θ_b as $V_i(x, z, a, S; \theta_b)$. Consider an alternative policy θ_a with value function $V_i(x, z, a, S; \theta_a)$. The alternative policy implies a change Δ_i in consumption equivalent units relative to the benchmark policy, where:

$$\Delta_i = \left(\frac{V_i(\theta_a)}{V_i(\theta_b)} \right)^{\frac{1}{(1-\gamma)\alpha_c}} - 1.$$

We compute aggregate welfare effects from a policy change by summing Δ_i across agents, calling the resulting aggregate welfare measure Δ. We also sum separately among owners, market renters, and RC renters, for different age groups, and for different income and wealth groups. In addition to computing steady state welfare effects, we can also compute transition paths that take into account the (slow) adjustment of the endogenous state variables to the new policy.

3 Calibration

We calibrate the model to the New York metropolitan area. Data sources and calculation details are described in Appendix B. Table 1 summarizes the chosen model parameters.

Geography The New York metro consists of 25 counties located in New York (12), New Jersey (12), and Pennsylvania (1). We assume that Manhattan (New York County) represents zone 1 and the other 24 counties make up zone 2.[22] The zones differ in the maximum buildable residential square footage, as captured in the model by $\overline{H}$. Appendix B describes detailed calculations on the relative size of Manhattan and the rest of the metro area, which imply that $\overline{H^1} = 0.0238 \times \overline{H^2}$. We then choose $\overline{H^2}$ such that the fraction of households living in zone 1 equals 10.5% of the total, as in the data. Since the model

[22] Alternative choices are to designate (i) New York City (five counties coinciding with the five boroughs of NYC) as zone 1 and the rest of the metro as zone 2, or (ii) Manhattan as zone 1 and the other four counties in New York City as zone 2. Both choices ignore that the dominant commuting pattern is from the rest of the metro area to Manhattan.

has no vacancies, we equate the number of households in the model with the number of occupied housing units in the data.

Demographics The model is calibrated so that one model period is equivalent to 4 years. Households enter the model at age 21, work until age 65, and retire with a pension after age 65. Survival probabilities $p(a)$ are calibrated to mortality data from the Census Bureau.[23,24] People above age 65 comprise 21% of the population above age 21 in the data and in the model.

Labor Income Since our model is an incomplete markets model in which housing affordability policies could act as an insurance device, it is important to realistically calibrate labor income risk and the progressivity of the tax and transfer system. Pre-tax labor income for agent i of age a is:

$$y_t^{i,a} = W_t n_t^i G^a z_t^i,$$

where the household takes wages as given and chooses labor supply n_t^i. Efficiency units of labor $G^a z_t^i$ consist of a deterministic component that depends on age (G^a) and a stochastic component z^i that captures idiosyncratic income risk.

The G^a function is chosen to enable the model to match both the evolution of the mean and the variance of earnings over the life cycle. We use data from ten waves of the Survey of Consumer Finance from 1983-2010 to estimate G^a. We refer the reader to appendix C for details.

The idiosyncratic productivity process z^i is chosen to both match labor income income inequality in the NYC data and to generate realistic persistence in earnings. We discretize z as a 4-state Markov chain. We set the annual income cutoffs at \$50,000, \$100,000, and \$200,000. This results in four income groups with average income of \$24,930, \$73,461, \$141,817, and \$343,693 in the NYC data. After scaling by NYC average income, we obtain the grid $Z = [0.255, 0.753, 1.453, 3.522]$. We set our productivity states equal to this grid.

The transition probability matrix is age-invariant but is allowed to depend on β type. Specifically, the expected duration of the highest productivity state is higher for the more patient agents. There are five unique parameters governing transition probabilities which

[23] To speed up computation, we assume that the probability of dying is zero before age 44. The observed probability is below 1% for each 4-year period before age 44. When the number of agents is not sufficiently large, a small probability of death induces idiosyncratic demographic risk, which leads to idiosyncratic variation in the wage. Smoothing out this idiosyncratic variation would require a very large number of agents (as opposed to when the probability of death is larger).

[24] We use mortality tables from 1960 rather than the latest available ones so as to generate the observed share of agents above age 65 in the current population.

are pinned down by five moments in the data. The four income groups have population shares in the data of 39.1%, 27.2%, 23.7%, and 10.0%, respectively. Since the shares sum to 1, that delivers three restrictions on the transition matrix. Matching the persistence of labor income to 0.9 delivers a fourth restriction. Finally, the dependence on β is calibrated to deliver the observed correlation between income and wealth in the SCF data. Appendix C contains the parameter values and further details.

Taxation We follow Heathcote, Storesletten, and Violante (2017) and choose an income tax schedule that captures the observed progressivity of the tax code in the U.S. in a parsimonious way:

$$T(y_t) = y - \lambda y^{1-\tau}$$

We set τ, which governs the progressivity, to 0.181, which is the value estimated by Heathcote et al. (2017). We set λ such that we balance the social security budget and generate a replacement rate of 56%. $RR = .56$ is the aggregate replacement rate observed in New York, then budget balance reads:

$$RR \int y_i di = \int T(y_i) di = \int y_i di - \lambda \int y_i^{1-\tau} di$$

Using the endogenous income distribution for NY in the model and the value for τ, this delivers $\lambda = 0.74$.

Retirement Income Retirement income is increasing in the household's last productivity level prior to retirement, but is capped for higher income levels. We use actual Social Security rules to estimate each productivity group's pension relative to the average pension. The resulting pension income states are $\psi^z = [0.520, 1.147, 1.436, 1.436]$. They are multiplied by average retirement income $\overline{\Psi} = 0.268$, which is endogenously determined in equation (9).

Commuting Cost Commuting has both a time cost and a financial cost. We choose the time cost to match the time spent commuting for the average New York metro area resident. This is 6.2 hours per week, according to Office of the Comptroller of New York City, itself based on Census and American Community Survey data. This time cost is the average of all commutes, including those within Manhattan. Since the model normalizes sets the commuting time within zone 1 to zero, we target the additional commuting time of zone 2 residents. The additional commuting time amounts to 25 minutes per trip for

Table 1: Calibration

Description	Parameter	Value	Target
Panel A: **Geography and Income**			
Productivity states	Z	[0.255 0.753 1.453 3.522]	Income inequality - see App. C
Transition prob.	$(p_{11}, p_{22}, p_{33}, p_{44}, p^H)$	(.933,.806,.780,.632,.238)	Pop. shares of income groups, persistence of income - see App. C
Tax progressivity	τ	0.181	Heathcothe et al. (16)
Tax burden	λ	0.74	Pension tax revenue
Average pension	$\overline{\Psi}$	0.268	Average replacement rate of 56%
Pension distr.	ψ^z	[0.520 1.147 1.436 1.436]	U.S. pension rules
Available space	$(\overline{H^1}, \overline{H^2})$	(0.169,7.10)	Max. residential buildable area, pop. share Manhattan
Time-Commuting cost	ϕ_T^2	0.037	Average commuting time NY
Financial-Commuting cost	ϕ_F^2	0.0079	Average commuting cost NY
Panel B: **Preferences**			
Risk aversion	γ	5	Standard value
Leisure weight	α_n	0.500	Average 42.8-hour workweek
Housing consumption weight	α_h	0.097	Avg. rent/avg. income of 29.3%
Time Preference (4yr)	(β^H, β^L)	(0.999,.80)	Average wealth/income 6.7 and wealth Gini 0.75 in SCF
Cons. externality threshold	$\underline{c}$	0.7	House-price/income ratio of zones in NY data
Extra utility zone 1 workers	χ^W	1.0025	Income ratio of zones in NY data
Extra utility zone 1 retirees	χ^R	1.1023	Fraction of retirees ratio of zones in NY data
Panel C: **Finance, Housing, Construction**			
Bond Price (4yr)	Q	0.914	Annual bond yield
Maximum residential LTV	θ_{res}	0.90	Mortgage underwriting standards
Maximum investment LTV	θ_{inv}	0.80	Mortgage underwriting standards
Property tax (4yr)	τ^P	0.0755	NY metro avg. property tax rate of 1.89% per year
Depreciation rate	δ	.0954	Residential depreciation rate of 2.39% per year
Panel D: **Production**			
Return to scale consumption sector	ρ_c	0.66	Labor income share of 2/3
Return to scale housing sector	ρ_h	0.4318	Housing supply elasticity for NY
Panel E: **Rent Control**			
Fraction rent control	(η^1, η^2)	(11%,6%)	Fraction of renters in RC of 16.9% and 10.4%
Rental discount	κ_1	50%	Observed rental discount
Income threshold for RC	κ_2	67%	Fraction of income Q1 hhs in RC
Rental share threshold for RC	κ_3	9.5%	Fraction of income Q3/Q4 hhs in RC

10 commuting trips per week. The 4.2 hours represent 3.7% of the 112 hours of weekly non-sleeping time. Hence, we set $\phi_T^2 = 0.037$.[25]

The financial cost of commuting is set at $\phi_F^2 = .0079$ in order for the model to match the 68% difference in market rents between Manhattan and the rest of the metro area. Market rents are calculated using Zillow's constant-quality rental index data, as detailed in Appendix B. This cost parameter implies that 2.1% of average working-age labor income is spent on commuting, or about $2,000 per household per year. This is a realistic cost for New York.[26]

We assume that retirees have time and financial commuting cost that are 1/3 of those of workers. We envision that retirees make fewer trips, travel at off-peak hours, and receive transportation discounts.

Preferences The functional form for the utility function is given in equation (1). We set risk aversion $\gamma = 5$, a standard value in the asset pricing literature.

Total non-sleep hours are normalized to one. We set $\alpha_n = 0.50$ to match the average workweek for New York residents. The observed average workweek is 42.8 hours or 38.2% of available time.

We set $\alpha_h = 0.097$ in order to match the ratio of average market rent to metro-wide average income. Income data discussed above and rental data from Zillow indicate that this ratio is 29.3% for New York in 2015. The value for $\alpha_c = 1 - \alpha_h - \alpha_n$ is implied.

We set $\beta^H = 0.999$ (0.999 per year) and $\beta^L = 0.80$ (0.945 per year). A 25% share of agents has β^H, the rest has β^L. This delivers an average β of 0.847, chosen to match the average wealth-income ratio which is 6.4 in the SCF data. The dispersion in betas delivers a wealth Gini coefficient of 0.88, close to the observed wealth Gini coefficient of 0.75 for the U.S.

For the taste-shifter for zone 1, we choose $\chi^W = 1.0025$, $\chi^R = 1.1023$, $\underline{c} = 0.70$. The latter number implies that 8.5% of the population would derive extra utility from living in Manhattan. We chose these three parameters to get our model to match the ratios of zone 1 relative to zone 2 variables: the relative fraction of retirees' of 0.91 (0.50 in model), a relative household income ratio of 1.42 (1.79 in model), and the relative price-to-income

[25]The 25 minute additional commute results from a 15 minute commute within Manhattan and a 40 minute commute from zone 2 to zone 1. With 10.5% of the population living in Manhattan, the average commuting time is 37.4 minutes per trip or 6.2 hours a week. This is exactly the observed average for the New York metro from Census data.

[26]In NYC, an unlimited subway pass costs $1,400 per year for a single person; many households have more than one commuter. Rail passes from the suburbs cost around $2400-3600 per year, depending on the railway station of departure. The cost of commuting by car are at least as high once the costs of owning, insuring, parking, and fueling a car are factored in.

ratio in zone 1 to zone 2 of 2.10 (1.40 in the model).

Housing We choose a price for the one-period (4-year) bond of $Q = 0.914$. This implies an annual, real yield-to-maturity of 2.27% on the four-year bond, which is the average observed in U.S. data.[27]

We set the maximum loan-to-value ratio (LTV) for the primary residence at 90% ($\theta_{res} = 0.9$), implying a 10% down payment requirement. The observed mean combined LTV ratio at origination for U.S. mortgages in the U.S. is 87.3% as of October 2016 according to the Urban Institute and has consistently been above 80% since the start of the data in 2001. The LTV for investment property is set at 80% ($\theta_{inv} = 0.8$), consistent with higher downpayment requirements for investment properties.

We set the property tax rate $\tau^P = 0.0755$ or 1.887% per year. This is the weighted average property tax rate for the states of New York (1.65%) and New Jersey (2.38%), taking into account that 1/3 of the NY metro area housing stock is located in New Jersey.[28]

We assume that property depreciates at 2.39% per year and set $\delta = 0.0954$. This is the average depreciation rate for privately-held residential property in the BEA Fixed Asset tables for the period 1972-2016.

Production and Construction We assume that the return to scale $\rho_c = 0.66$, which implies a labor share of 66%, consistent with the data.

For the housing sector, we set $\rho_h = 0.432$ in order to match the housing supply elasticity, given the other parameters. The long-run housing supply elasticity in the model is given in appendix D. Saiz (2010) reports a housing supply elasticity for the New York metro area of 0.76. The model delivers 0.756. As an aside, the housing supply elasticity is much lower in zone 1 (0.019) than in zone 2 (0.744), because in zone 1 the housing stock is much closer to $\overline{H}$ (23.5% from the constraint) than in zone 2 (86.4% from the constraint). Since the housing stock of the metro area is concentrated in zone 2, the city-wide elasticity is dominated by that in zone 2.

Rent Control Rent regulation plays a major role in the New York housing market. According to the New York City Housing and Vacancy Survey and county-level data on af-

[27]The five-year constant maturity Treasury yield averaged 5.56% while the one-year CMT Treasury note averaged 4.91% between 1953.Q2 and 2017.Q3, the longest available time series. The linearly interpolated four-year yield to maturity is 5.40%. Inflation is 3.13% on average over the same period, so that the real yield is 2.27%.

[28]For New York City, the property-value weighted average residential property tax rate paid as a fraction of estimated market value is 1.65%. No comparably precise data are available for the other 20 counties in zone 2.

fordable housing for the New York metro area counties outside of New York City, 16.9% of the renters in zone 1 and 10.4% of the renters in zone 2 live in rent-regulated units. The average rent in rent-regulated properties is 49.9% below that in all other rentals. Appendix B.5 contains a detailed description of our rent regulation data and definitions. We set the square feet of housing devoted to RC units to $\eta^1 = 5\%$ and $\eta^2 = 6\%$ to match the fraction of renting households that are in rent controlled units in each zone.

We set $\kappa_1 = 0.50$ to the observed rent discount. It follows that $\kappa_4^1 = 0.945$ and $\kappa_4^2 = 0.97$, so that landlords earn 5.5% lower average rents in zone 1 and 3% lower average rents in zone 2 than they would in an unregulated market.

We set the income qualification threshold to a fraction $\kappa_2 = 67.3\%$ of AMI. The target for this parameter is the fraction of households who are in RC housing, conditional on being in the first quartile of the city-wide household income distribution. For the RC size threshold, we set $\kappa_3 = 9.54\%$ to match the fraction of households who are in RC, conditional on being in the top half of the income distribution. High-income agents in the model want to live in a house that is larger than the maximum allowed size under RC.

4 Baseline Model Results

We start by discussing the implications of the baseline model for the spatial distribution of population, housing, income, and wealth. We also discuss house prices and rents for the city as a whole and for the two zones. Then we look at the model's implications for income, wealth, and home ownership over the life-cycle.

4.1 Demographics, Income, and Wealth

Demographics By virtue of the calibration, the model matches the population share in each of the two zones. Manhattan contains 10.5% of the population of the metro area (NYC). The ratio of households in zone 1 to zone 2 is 11.8%; see first row of Table 2.

The average age of NYC residents is 47 years old in the model, and lower in Manhattan (41) than the rest of the metro (48). Younger households are in the working phase of life, when proximity to work is valuable. Younger agents tend to have lower income, which increases the importance of not having to bear the financial cost of commuting. However, they may not be able to afford the high rents in Manhattan until later in their life-cycle when income is higher.

People aged 65 or older represent 14% of the population over 21 in Manhattan, a lower share than in the rest of the metro. Since retirees do not work, they do not need to live close to work. Furthermore, they enjoy lower time and financial costs of commuting, further weakening their incentives to live in Manhattan. On the other hand, retirees tend to be wealthier making living in Manhattan financially feasible. The taste-shifter for high-consumption retirees, χ^R, is instrumental in getting enough retirees to live in Manhattan.

Housing Units In the data, the typical housing unit is much smaller in Manhattan than in the rest of the metro area. We back out the typical house size (in square feet) in each county from the median house value and the median house value per square foot, using 2015 year-end values. We obtain an average housing unit size of 1,400 sqft on average, but only 826 sqft in Manhattan. The ratio of zone 1 to zone 2 is 0.56. In the model, households freely choose their housing size. The model generates a ratio of house size in zone 1 to zone 2 of 0.96.

Multiplying the number of households, which equal the number of housing units, by the average size per housing unit, we find that Manhattan contains 6.2% of the overall housing stock of the metro compared to 10.5% of the population. The current housing stock represents 34% of the maximum buildable residential square feet in zone 1 but only 12% of the maximum in zone 2.

Income Average income in the metro area is slightly higher than in the data. The ratio of average income in zone 1 to zone 2 is also higher than in the data (row 5 of Table 2). Productive working-age households have a high opportunity cost of time and prefer to live close to work. Partly offsetting this force is the fact that the financial cost of commuting weighs more heavily for low-income households, which pushes them towards living in zone 1. Rent control permits lower income households to live in zone 1 and lowers the average income in zone 1. On balance, we require a minor taste-shifter χ^W to affect the ratio of income in zone 1 to zone 2.

The top left panel of Figure 1 shows household income over the life-cycle, measured as pre-tax labor income in working age and social security income in retirement. We plot average income in the bottom 33% of the income distribution, in the middle of the income distribution (33-67%), in the 67-90th percentiles, and in the top-10% of the distribution. Labor income has a hump-shaped profile over the life-cycle inherited from the deterministic productivity process. The model generates a large amount of income inequality at every age. The model's overall income Gini is 0.56; it was 0.503 in 2015 in the NY metro data. The top right panel of Figure 1 plots average income profiles by zone. Average income is higher in Manhattan throughout the life cycle, for reasons explained above.

Table 2: New York Metro Data Targets and Model Fit

		NY Data		NY Model	
		metro	**ratio zone 1/zone 2**	**metro**	**ratio zone 1/zone 2**
1	Households (thousands)	7,124.9	11.8%	7,124.9	11.5%
2	Average hh age	38.0	1.03	47.2	0.87
3	People over 65 as % over 21	19.3	0.91	21.3	0.60
4	Average house size (sqft)	1,399	0.56	–	0.96
5	Pre-tax income per hh ($)	97,577	1.42	128,010	2.33
6	Home ownership rate (%)	50.8	0.44	66.0	1.30
7	Median mkt price per unit ($)	495,694	2.85	358,910	1.52
8	Median mkt price per sqft ($)	354.3	5.06	–	1.62
9	Median mkt rent per unit (monthly $)	2,381	1.66	1,796	1.46
10	Median mkt price/median mkt rent (annual)	17.35	1.72	16.66	1.04
11	Mkt price/avg. income (annual)	5.08	2.01	3.90	0.61
12	Avg. rent/avg. income (%)	29.3	1.17	22.7	1.11
13	Avg. rent/income ratio for renters (%)	42.0	0.93	48.2	0.85
14	Rent burdened (%)	53.0	0.89	23.6	0.03
15	Severely rent burdened (%)	31.0	0.88	34.5	1.02
16	% Rent regulated	11.5	1.63	10.5	1.86

Notes: Columns 2-3 report the values for the data of the variables listed in the first column. Data sources and construction are described in detail in Appendix B. Column 3 reports the ratio of the zone 1 value to the zone 2 value in the data. Column 5 reports the same ratio in the model.

Wealth The model makes predictions for average wealth, the distribution of wealth across households, as well as how that wealth is spatially distributed. Wealth inequality is high, with a wealth Gini coefficient of 0.88, close to the data.

The probability of receiving a bequest equals the number of households between ages 21 and 65 divided by the number of dead households. It is equal to 10% over each 4-year period, and identical for β^H and β^L household types. Under our calibration, about 1.2% of wealth is bequeathed each year. This number matches the data. Some agents with low productivity who receive a bequest decide not to work and choose to live in zone 2.

The model predicts a ratio of average wealth in zone 1 to zone 2 of 3.79. While this number is not observable in the data, it seems reasonable to expect wealth inequality across parts of the city that mirrors income inequality.

The middle left panel of Figure 1 plots household wealth over the life cycle at the same *income* percentiles. It shows that the model endogenously generates substantial wealth accumulation for the average New York resident and a very large amount of wealth inequality between income groups, which grows with age. The middle right panel shows how average wealth evolves differentially across zones.

Figure 1: Income, Wealth, and Home Ownership over the Life-Cycle

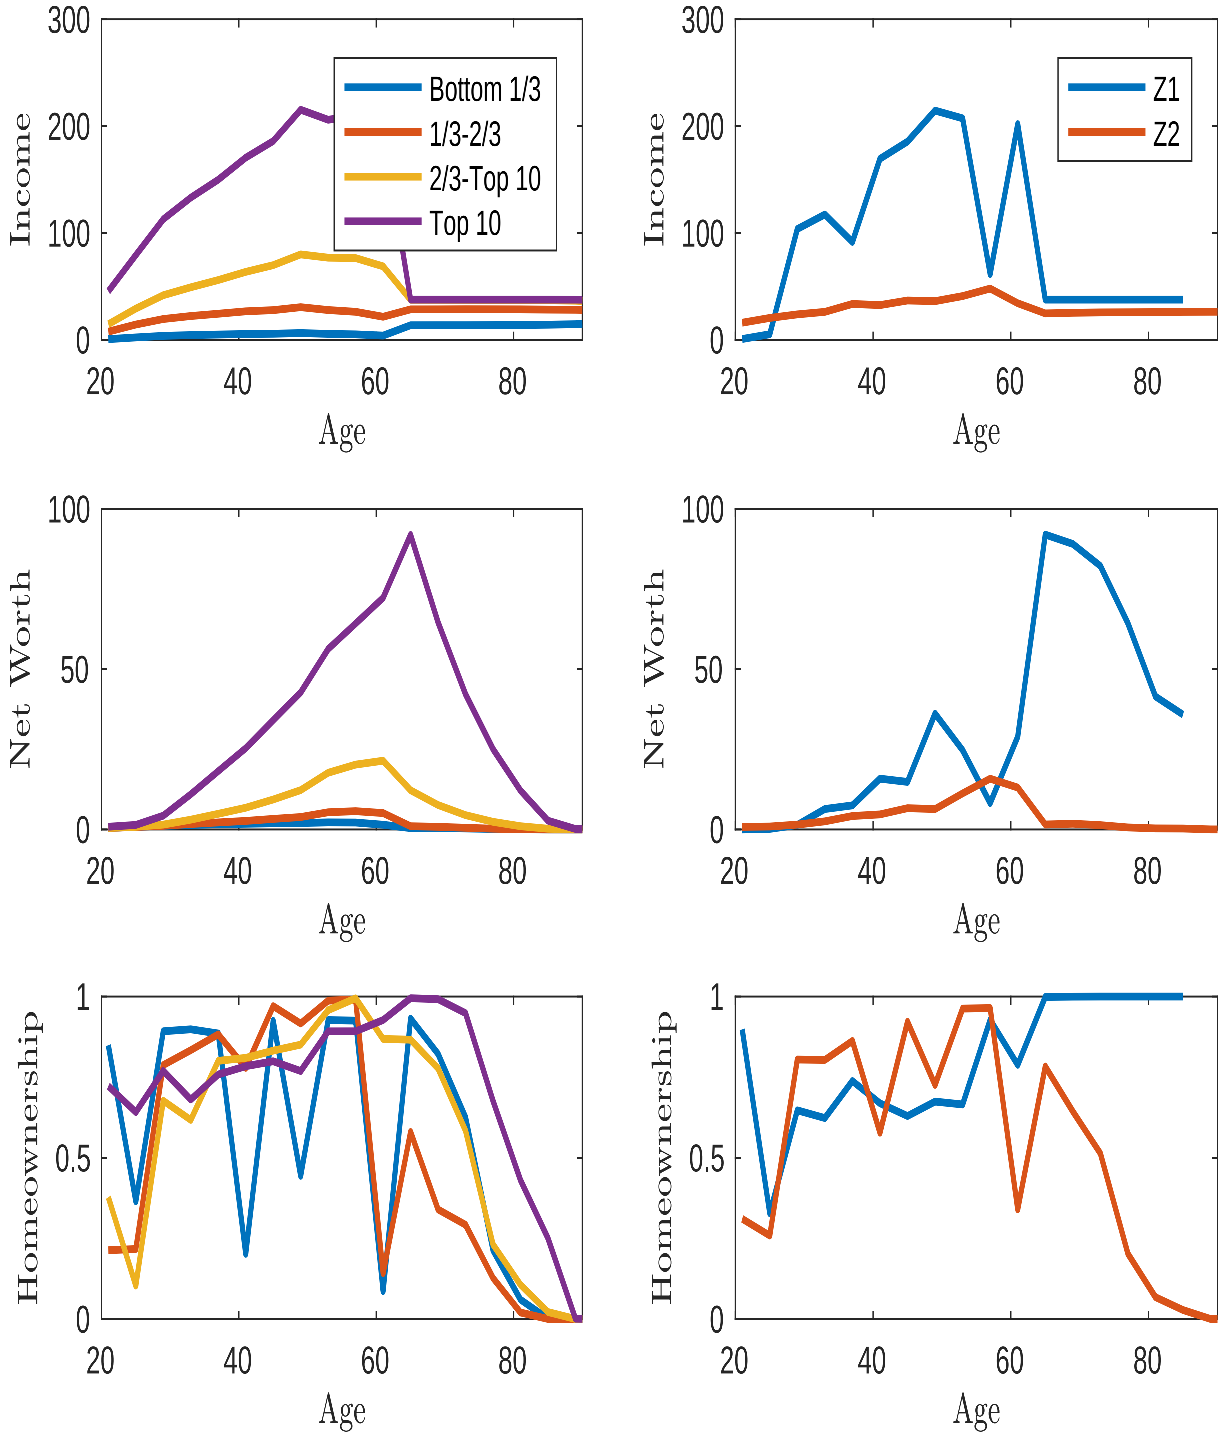

4.2 Home Ownership, House Prices, and Rents

Next, we explore the model's predictions for home ownership, house prices, and rents. The model manages to drive a large wedge between house prices, rents, income, and home ownership rates between zones 1 and 2 for realistic commuting costs.

Home Ownership The bottom left panel of Figure 1 plots the home ownership rate over the life-cycle. The bottom right panel of the figure plots the average home ownership rate profiles by zone.

The model generates a home ownership rate of 66%, somewhat above the 50.8% in the New York metro. Row 6 of Table 2 shows that the home ownership rate in Manhattan, at 23.5%, is far below that in the rest of the metro area, at 54.1%. The ratio of these two numbers is 0.43. The model is unable to generate a lower home ownership in zone 1 than in zone 2. It gets the home ownership rate in zone 2 about right, but overstates the home ownership rate in Manhattan. The low home ownership rate in Manhattan is consistent with its higher high price-rent ratio (see below).

Market Prices and Rents Turning to rents and house prices, row 7 of Table 2 shows the median price per housing unit, row 8 the median price per square foot (the ratio of rows 8 and 2), row 9 the median rent per unit, and row 10 the price-rent ratio. In the data, we use the Zillow home value index (ZHVI) to measure the median price of owner-occupied units, the Zillow median home value per square foot, and the Zillow rental index (ZRI) for the median rent. These indices are available for each county in the New York metro, and we use the year-end 2015 values. Zillow excludes non-arms' length transactions and rent-controlled rentals. To aggregate across the 24 counties in zone 2, we calculate the median price as the weighted average of the median prices in each county, where the weights are the shares of owner-occupied units. Similarly, for the median rent of zone 2, we average median rents of the 24 counties using renter-occupied unit shares as weights. Zillow uses a machine-learning algorithm that ensures that the ZHVI and ZRI pertain to the same, typical, constant-quality unit, in a particular geography. The ratio of the ZHVI to the ZRI in a county, is the price-rent ratio.

To ensure consistency with the empirical procedure, we calculate the median house size in each zone including both owner- and renter-occupied units (but excluding RC units) in the model. Call these $\bar{h}^{\ell}$. We form the median price per unit as the product of the market price times the typical unit size $P^{\ell}\bar{h}^{\ell}$. The market rent is $R^{\ell}\bar{h}^{\ell}$. The price-rent ratio is simply $P^{\ell}\bar{h}^{\ell}/R^{\ell}\bar{h}^{\ell} = P^{\ell}/R^{\ell}$.

The median house value in the NY metro area is $495,694 in the data compared to $359,910 in the model. The median is $1,182,500 in Manhattan and $414,742 outside Manhattan in the data, a ratio of 2.85. This 2.85 house value ratio is the product of a house size ratio of 0.56 and a price per sqft ratio of 5.06. The model generates a ratio of prices of 1.52, a ratio of sizes of 0.96, and a ratio of price per sqft of 1.62.

The Zillow data indicate a monthly rent on a typical market-rate unit of $3,697 in Manhattan and $2,226 in zone 2; their ratio is 1.66. The model targets and matches this ratio through its choice of the working-age utility shifter χ^W. The predicted market rent level for the entire metropolitan area is $1,796, compared to the observed metro average of $2,381.

The land scarcity, proximity to jobs, and proximity to amenities (for those who are above the threshold) are the fundamental reasons why the model generates higher housing demand in Manhattan. Because of the highly inelastic housing supply in Manhattan, this translates into much higher house prices and rents.

The model comes close to capturing the metro-wide price/rent ratio level of 17.35 (row 10). A simple user cost model would imply a steady state 4-year price-rent ratio of $(1 - Q \times (1 - \delta - \tau^P))^{-1} = 4.13$ or 16.5 when rent is expressed as an annual number. The model has borrowing constraints so that housing has an additional collateral value component which increases its price. There also is a small housing risk premium.[29] In the data, the price-rent ratio in Manhattan is 26.65, or 1.72 times the 15.52 value in zone 2. In the model that ratio is 1.04. In other words, the model does not generate enough spatial variation in price to rent ratios. Manhattan may be desirable for reasons of proximity to work or to amenities, but those benefits accrue equally to renters and to owners.

Owners in zone 1 live on average in slightly smaller units than market renters in zone 1. The average value of owners' homes is slightly below the constant-size median home value reported in the table. Average rent paid by renters would be above the reported constant-size rent. The ratio of the two would exceed the reported price-rent ratio (we obtain a value of 17.66), and would differ across zones.

4.3 Housing Affordability and Rent Control

Price-Income and Rent-Income Row 11 of Table 2 reports the ratio of the median value of owner-occupied housing to average income in each zone. Income is averaged among all residents in a zone. The median home price to the average income is an often-used

[29] The model has no aggregate risk and risk aversion is only 5. To generate meaningful variation in housing risk premia would require, for example, changes in mortgage lending standards as in Favilukis et al. (2017).

metric of housing affordability. In the NY metro data, the median house costs 5.08 times average income. Price-income is 8.90 in Manhattan compared to 4.44 outside Manhattan, a ratio of 2.01. The model generates a ratio of 0.61. This is a consequence of the model imperfectly matching the average income ratio across zones, and of not generating enough spatial variation in median house values.

Next, we turn to renters. Row 12 reports average rent paid by market renters divided by average income of all residents in a zone. This moments was the target for the housing preference parameter α^h, and the model is slightly below 29.3% target. Computing this housing expenditure ratio separately for each zone, we obtain a value of 33.4% for Manhattan and a ratio of zone 1 to zone 2 of 1.17 indicating that rent differences are more pronounced than income differences. The model matches this pattern, with a ratio of 1.11.

To get at the individual rent burden, we compute three additional moments reported in row 13-15 of Table 2.[30] The first one computes household-level rent to income ratios for renters with positive income, caps the ratios at 101%, and takes the average across households. This average share of income spent on rent is 40% in Manhattan and 43% in New York City ex-Manhattan. The second one computes the fraction of renters whose rent is between 30% and 50% of income. These households are known as rent-burdened. The third measure computes the fraction of renters whose rent exceeds 50% of income, as well as households with zero income. This group is known as severely rent-burdened. In the Manhattan data, 48% of renters are rent burdened and another 28% are severely rent burdened. Outside of Manhattan, 54% of renters are rent-burdened and an additional 32% severely rent-burdened. The model exactly matches the average rent to income ratio and the fraction of severely rent-burdened households, while understating the fraction of moderately rent-burdened households.

Rent Control By virtue of the calibration, the model generates the right share of rent-controlled households in renter-occupied housing in each zone (row 16 of Table 2). The calibration also targeted the distribution of rent controlled households into the various income quartiles. In the model, most high- and middle-income households who win the lottery turn down rent-regulated housing. The maximum rent or maximum income restrictions are too unappealing from a utility perspective (recall labor supply is endogenous). Most low-income households accept the lottery if they win. Their labor supply choice is unaltered since their optimal choice of hours implies an income that is below the

[30] We only have household-level data for the five counties in New York City, not for the zone 2 counties outside of NYC. The implicit assumption is that the values in these counties are the same as in the four outer boroughs of NYC. One caveat is that the NYHVS rent levels are substantially below the Zillow rents. The ratio of zone 1 to zone 2 in both data sets are nearly identical.

Figure 2: Distribution of rent-controlled agents by age, and income and net worth quartiles.

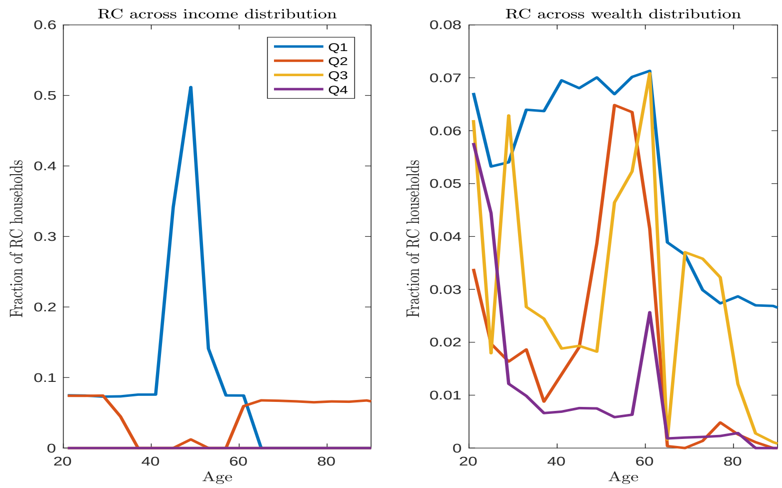

threshold. Some lower-middle income households who win the lottery end up reducing their hours in order to meet the income criterion. This has adverse implications for the city-wide labor supply and production. A second distortion is that low-income households who win the lottery and are unconstrained by the maximum rent demand more housing than they would under market conditions.

Rent control acts as an insurance device in our incomplete markets model. If it is difficult for a low-productivity household to get into the RC system, then the value of the insurance is low. The model predicts that, conditional on being in the lowest income quartile, the chance of winning and accepting the RC lottery (in either zone 1 or zone 2) is 9.1%, while it is 5.2% for the second quartile, and decreases to zero for the third and fourth quartiles.

Figure 2 plots the model's fraction of households that are in RC for the four income quartiles against age in the left panel, and for the four wealth quartiles in the right panel. As implied by the acceptance probabilities of richer households, no households are rent-controlled in the third and fourth quartile of the income distribution. The fraction of RC households is approximately constant over age in the first two quartiles, but peaks before retirement age, when the least productive households reach the peak of their age-dependent productivity and choose not to work to avoid facing progressive taxation, thus falling under the income cutoff for RC. Across the wealth distribution, younger and poorer households tend to be more rent-controlled than older and richer ones.

5 Affordability Policies

This section studies various policy experiments that change key levers associated with the rent control system and with zoning laws. Table 3 summarizes how the main moments of the model change under the six policy experiments. The first column reports the benchmark model, while the other columns report the percentage change in the moments in deviation from the benchmark. The first set of moments describes changes in housing affordability, the second set describes changes in aggregates across the two zones, and the third set describes how welfare changes are distributed across wealth levels. Figure 3 plots the associated welfare changes for the first four policies that focus on RC, across the age distribution and tenure statuses. Figure 4 does the same for the last two policies that modify the spatial equilibrium of the city. Welfare changes are measured as permanent changes in consumption equivalent units.

Table 3: Main moments of the model under affordability policies that modify features of the RC system and the spatial allocation of housing.

		Benchmark	RC share	RC discount	Inc. cutoff	RC housing size	All RC in Z2	Zoning Z1
1	Avg. rent to income Z1	1.04	-3.61%	-8.65%	4.78%	7.50%	-8.63%	35.76%
2	Avg. rent to income Z2	0.59	-9.84%	-3.74%	7.04%	7.69%	2.86%	-4.17%
3	Frac. RC renters	0.11	-47.71%	47.54%	0.47%	84.74%	6.85%	5.81%
4	Frac. RC in income Q1	0.64	0.45%	56.76%	29.19%	57.80%	-1.71%	0.16%
5	Frac. sev. rent-burdened	0.12	-19.21%	1.79%	7.14%	22.13%	0.63%	-29.25%
6	Avg. size RC unit	0.70	0.12%	-32.83%	-0.18%	-49.39%	3.24%	0.82%
7	Avg. size Z1 unit	1.36	-2.54%	-32.74%	-0.64%	5.08%	0.47%	-23.20%
8	Avg. size Z2 unit	1.43	0.49%	-32.86%	0.07%	-0.69%	-0.20%	4.46%
9	Housing stock Z1	0.14	0.04%	0.17%	0.01%	0.11%	-0.14%	8.62%
10	Housing stock Z2	1.28	0.23%	0.17%	0.00%	-0.14%	-0.15%	-0.16%
11	Rent/sqft Z1	0.09	-0.10%	-0.89%	0.20%	-1.50%	-0.13%	-1.13%
12	Rent/sqft Z2	0.06	-0.09%	-0.71%	0.16%	-1.18%	-0.10%	-1.02%
13	Price/sqft Z1	0.39	-0.02%	-0.21%	0.05%	-0.35%	-0.03%	-0.26%
14	Price/sqft Z2	0.24	-0.02%	-0.17%	0.04%	-0.29%	-0.02%	-0.25%
15	Homeownership Z1	0.84	1.87%	-1.24%	-0.10%	-3.79%	2.97%	1.93%
16	Homeownership Z2	0.64	2.76%	-0.34%	0.24%	-3.34%	-0.61%	4.48%
17	Frac. pop. Z1	0.10	2.29%	0.11%	0.67%	-4.82%	-0.43%	38.45%
18	Avg. income Z1	0.72	-3.95%	-1.65%	-0.75%	6.09%	0.76%	-35.32%
19	Avg. income Z2	0.31	0.54%	0.17%	-0.01%	-0.91%	-0.18%	6.53%
20	Avg. productivity Z1	1.88	-3.63%	-1.92%	-0.70%	4.75%	0.50%	-31.23%
21	Avg. productivity Z2	0.92	0.39%	0.16%	-0.01%	-0.62%	-0.11%	5.03%
22	Output consumption	1671.41	-0.10%	-0.09%	-0.14%	-0.06%	-0.07%	-0.18%
23	Output construction Z1	20.63	0.07%	-1.19%	-0.07%	-1.49%	1.92%	9.16%
24	Output construction Z2	116.94	0.66%	0.42%	0.02%	-0.46%	-0.30%	-0.37%
25	Welfare change NW Q1	–	-0.06%	-0.23%	-0.06%	-0.19%	0.03%	0.97%
26	Welfare change NW Q4	–	0.20%	0.12%	-0.01%	-0.10%	-0.03%	0.60%

Notes: Column “Benchmark” reports values of the moments for the benchmark model. Columns “RC share” to “Zoning Z1” report percentage changes of the moments in the policy experiments relative to the benchmark. Rows 1-8 report housing affordability moments , rows 9-24 aggregate moments across the two zones, and rows 25-26 welfare moments.

5.1 Reducing the Fraction of Rent Controlled Housing

In the first experiment, we reduce the fraction of rent controlled housing (in square feet) in each zone. The new η^{ℓ} values are half of the benchmark values. The fraction of *households* in RC falls by 50%. With the fall in the probability of winning the RC lottery comes an increase in the probability of accepting RC conditional on winning, especially in zone 1 (+4%). The distribution of RC units across the quartiles of the income distribution remains essentially the same as in the benchmark economy. Of those in RC, 63.7% are in the first quartile of the income distribution and 36.3% in the second quartile.

The reduction in RC results in lower rents: -0.10% in Manhattan and -0.09% in zone 2. Lower rents attract more households to Manhattan. The population share of Manhattan rises by 2.3%. Because the housing stock in zone 1 does not change much (+0.04%), accommodating more households induces the size of the average housing unit in zone 1 to shrink (-2.5%). While the typical RC unit has about the same size as in the benchmark, market rental units become 2.6% smaller in zone 1. The housing stock in zone 2 increases (+0.23%), as does the economy-wide housing stock, since developers face fewer distortions from rent control (κ_4 is closer to 1). This observation is consistent with the empirical literature, which finds that increased incentives of landlords to renovate their properties and of developers to invest in new construction generate a modest housing boom in decontrolled areas (Autor et al. (2014), Diamond et al. (2017)). Lower rents and a higher housing stock imply that the now fewer zone 2 households each live in larger housing units (+0.5%). Market rental units are even 2.5% larger on average in zone 2 compared to the benchmark economy.

Housing affordability improves. The fraction of severely rent burdened renters falls by 19% economy-wide. Average rent-income ratios fall for all four quartiles of the income distribution, with much larger reductions in the bottom half than in the top half of the income distribution. The improvements in affordability are stronger in zone 2 where rent to labor income ratios fall by 9.8%. They only fall by 3.6% in zone 1. This happens because some middle-income households move to zone 1 where they pay high rents. With less space constrained by RC, more households choose to live in Manhattan, and the improvements in affordability show up mostly in zone 2 statistics.

Median price to rent ratios fall in the entire city, resulting in higher home ownership rates in zone 1 (+1.9%) and in zone 2 (+2.8%). Most of the square footage freed up by the reduction in the RC share is absorbed by owner-occupied housing rather than by market rentals. This is also consistent with the conversion of previously RC units to renter-occupied units and condominiums observed in practice. Per capita wealth is lower in each zone because of lower house prices, but more so in zone 1 than in zone 2 so that the wealth gap between zones shrinks. The price-income ratio falls by 3.5% metro-wide, in a further sign of improved housing affordability and consistent with the increase in home ownership.

Relative to the benchmark, medium-productivity households migrate from zone 2 to zone 1. This is because commuting costs and amenities make living in zone 1 more desirable, and now rents and house prices are lower in zone 1. The productivity and income gaps between zones shrink. Per capita income in zone 1 is 4.0% lower than in the benchmark while per capita income in zone 2 rises by 0.5%. This reflects both smaller differences in productivity and in hours worked across zones. Income inequality (Gini coefficient) in the metro area falls by 0.2%. Income inequality within Manhattan rises by 2% as there is a greater diversity of productivity levels present. Thus reducing RC symmetrically across zones has the asymmetric effect of attracting previously crowded-out, poorer households to the more desirable part of the city[31]. If Manhattan possesses extreme-skill complementarity, as many large cities in the U.S. (Eeckhout et al. (2014)), increased economic diversity would further contribute to an increase in social welfare. Here, we do not model extreme-skill complementarity, and thus interpret our results as a lower bound on the increase in social welfare.

Turning to the macro aggregates, the output of the construction sector rises, especially in zone 2, while the output of the non-housing sector falls. Overall output falls slightly (-0.05%). This occurs despite a slight rise in the total hours worked (+0.28%) because of the stronger decreasing returns to scale in construction, which induces the aggregate productivity of labor to endogenously decrease. In spite of the construction booms in both zones, the misallocation of labor across sectors results in lower output. Because of the larger population share of Manhattan, total commuting time falls by 0.34%. Total leisure hours fall since the increase in hours worked exceeds the time saved commuting less.

The welfare effect of the policy sharply differs across the tenure distribution (Figure 3). Reducing the fraction of RC housing naturally hurts RC renters whose welfare falls by 1.43%. Market renters gain by 0.06% as housing affordability improves. Owners also gain from the policy change by 0.13%. Taking into account that the shares of these three groups in the population change in response to the policy, aggregate welfare increases by 0.052%. As partly implied by the effects of the policy across the tenure distribution, the latter is regressive: welfare gains are increasing in income and in wealth. In addition,

[31]This is the symmetric case of what happened when RC in Cambridge was removed, according to Autor et al. (2014). There, richer households flew into units previously occupied by poorer RC tenants.

gains and losses are amplified with age: before retirement, as older workers face more volatile idiosyncratic income profiles; and after retirement, when they lose the subsidy associated with the progressive tax and transfer schedule and just earn a fraction of their last labor income. More generally, policies scaling back RC programs generate a wedge between poor retirees living in RC units and rich retired homeowners.

5.2 Reducing the Rent Subsidy for RC Housing

In the second experiment, we reduce the rental discount that RC households enjoy. Specifically, we lower the rent discount parameter κ_1 from 50% to 25%, while keeping the share of square footage that goes to RC housing unchanged. Interestingly, this policy change results in about 46% more households in RC units in zone 1 and 49% more in zone 2. The fraction of all households (renters) in RC in the economy is 5.3% (15.5%) compared to 3.6% (10.5%) in the benchmark. The average RC unit is 33% smaller than before, reflecting the smaller discount.

RC units become less attractive for the average household, but more attractive for households at the bottom of the income distribution. Indeed, this policy is very effective at concentrating RC among the lowest income households. Among the households that are in RC, nearly all (99.4%) are in the bottom quartile. In the benchmark economy, we have 63.4% in quartile 1 and 36.6% in quartile 2. Among all households in income quartile 1, 21% now have RC while only 9% did in the benchmark. This result may seem surprising. One may have conjectured that reducing the rent subsidy would crowd out lower income households from RC. Instead, RC units become endogenously smaller to the point that households with incomes in the second quartile of the income distribution overwhelmingly turn down the RC unit conditional on winning the lottery. Those were trading-off the RC discount versus their apartment size. With a smaller discount, they prefer working more hours and being able to afford a larger housing unit on the private market.

The smaller RC distortion stimulates development. Across our four RC policies, the decrease in price distortions is the change that most stimulates construction in the whole city, because it directly restores developers' incentives to build, rather than relying on general equilibrium effects alone. Because RC units represented a relatively small share of new construction in the benchmark, the distortions faced by developers were modest in the first place. Therefore, the housing stock increases by 0.17% in each zone. Market rents fall by 0.89% in zone 1 and 0.71% in zone 2. Despite the additional housing supply available and the increase of households in RC, the population share of Manhattan stays constant. The reason is that the average market rental unit is 1% bigger.

Being more targeted on low income households, by inducing these households to choose smaller RC units, and by lowering overall market rental rates per square foot, this policy improves housing affordability. Households in the first income quartile now spend 9% less of their income on rent. This is despite the smaller subsidy rate. However, affordability deteriorates for households in income quartile 2 who lose RC and see their rent-income ratio rise by 7%. Renters in the top half of the income distribution spend about the same on rents as before, in part because they now live in larger units. Thus the fraction of households that is rent burdened or severely rent burdened does not fall.

With more RC households in Manhattan, per capita productivity in Manhattan falls by 1.9% and income falls by 1.65%. The latter reduction is smaller than the former because hours worked increases and offset it.

The home ownership rate falls, especially in Manhattan, which is consistent with the rise in the price to income and the price to rent ratios. Per capita wealth falls in each zone.

Overall output in the economy falls slightly (-0.07%), as an increase in construction output in zone 2 is offset by a reduction in non-housing production.

Aggregate welfare falls by 0.096%. RC renters' welfare falls by 2.74%, while market renters (+0.0043%) and owners (+0.0003%) gain very little. This policy is also regressive: there are welfare losses for households in the first and second income quartiles and gains for those in the third and fourth quartiles. The welfare gains are also increasing in wealth, only benefitting the top quartile of the wealth distribution. Figure 3 describes the distribution of welfare losses across ages and tenure statuses.

5.3 Reducing the Income Threshold for RC Housing

In the third experiment, we tighten the income requirements to qualify for rent control, governed by the parameter κ_2. Households must make less than 33.5% of AMI to qualify, compared to 67% in the benchmark economy. The fraction of households in RC in this policy experiment remains unchanged, but now RC households are more concentrated in the bottom quartile of the income distribution (81.9% of RC households) and less in the second quartile, compared to the benchmark. Surprisingly, the previous experiment was more effective at allocating RC units to the lowest income households than directly conditioning on income. Given the more stringent income threshold, households in the second quartile of the income distribution who win the RC lottery accept with only 33% probability, compared to 88% in the benchmark model. For those households, decreasing their hours to qualify for RC at the expense of their consumption of non-durable goods is not worth it. In the presence of a progressive tax and transfer schedule, this trade-off goes against entering RC for all but the poorest households. Overall, housing affordability changes little.

The policy change results in 0.7% more households living in Manhattan, on average living in housing units that are 0.6% smaller. Average income in zone 1 is 0.75% lower than in the benchmark as the average productivity of zone 1 workers falls (-0.7%). Output in the economy falls by 0.13%, driven by a 0.14% decline in output in the non-housing sector. Hours worked falls by 0.11%. Overall, this policy is mildly recessionary, and causes a decrease in production in all three sectors, which goes hand in hand with a decrease in average wealth and income in both zones.

As a result, aggregate welfare falls by 0.035%. RC renters' welfare falls by 0.67% while market renters (-0.018%) and owners (-0.009%) lose modestly (Figure 3). Renters lose because rents increase slightly (+0.17%), reflecting the increased demand for market rentals from middle-income households previously in RC. The policy fails to stimulate housing development; the housing stock is essentially unchanged.

5.4 Reducing the Maximum Size of RC Housing

The fourth experiment lowers the maximum size of a rent controlled unit, governed by the parameter κ_3. As in the other experiments, we lower it by 50%. Effectively, this implies that RC units become 50% smaller but the overall share of housing (square footage) devoted to RC remains the same. We can think of this policy as allowing for "micro units," as advocated by several policy institutes[32]. The fraction of all households (renters) in RC increases by 97% (85%) from 3.6% (10.5%) to 7.0% (19.4%). The probability of winning the RC lottery more than quadruples. However, conditional on winning, a lot more households turn down the RC unit because of its small size. This policy manages to concentrate 100% of the RC units in the hands of households in the first quartile of the income distribution; 28% of the households in that income quartile live in a RC unit. This illustrates another trade-off at work in social welfare: between fewer insiders previously in RC who are forced to downsize; and previously rent-burdened outsiders who lived in market rentals, who flow into RC once more, smaller units are available.

The policy has substantial effects on the non-regulated rental market. Rents per foot fall by 1.5% in Manhattan and by 1.2% in zone 2. The housing stock increases by 0.11% in Manhattan but falls by 0.14% in the rest of the metro. The average unit size in zone

[32]For instance the NYU Furman center (2018). New York City currently disincentivizes the development of small apartment units for a variety of regulatory reasons.

1 increases by a substantial 5.1% in zone 1 while falling by 0.7% in zone 2. The small increasing in the housing stock in Manhattan is more than offset by the large increase in the typical unit size so that the fraction of households in Manhattan falls by 4.8%. The remaining population in Manhattan is slightly older (+2%), much more productive (+4.7%), works more hours (+4.7%), and has much higher income (+6.0%) than in the benchmark economy. It is also wealthier (+4%). The opposite is true in zone 2. In other words, this policy creates more income and wealth inequality between zones by increasing the concentration of high-productivity workers in zone 1.

House prices fall but by less than rents, so that price-rent ratios rise. The home ownership rate falls in both zones. With larger units and the same total square footage dedicated to owner-occupied housing, there must be fewer owners in equilibrium.

Of all RC experiments, this policy has the largest negative effect on aggregate output (-0.10%), with all three sectors contributing to the decline. Total hours worked shrink (-0.58%). A larger share of households has access to heavily subsidized housing, which reduces their incentives to work. Total commuting hours increase by 0.67%.

It also has the largest negative effect on social welfare, which falls by 0.182%. This is because it is the policy that most directly affects the quantity of housing services consumed by RC households, inducing a first-order loss in their instantaneous utility. Indeed, RC renters lose 3.04%, and lose more across all age bins than in other experiments (Figure 3). Owners lose 0.11%, and market renters lose 0.0004%. The welfare losses occur in each quartile of the income distribution but become smaller for higher income levels.

5.5 Spatial Allocation of RC Housing

In the fifth experiment, we shift all rent controlled units from Manhattan to zone 2, keeping the overall fractions of households in rent control the same. The fraction of households in RC in zone 2 increases from 3.6% to 4.3%.

The left panel of Figure 4 summarizes the welfare effects of the policy. The latter has only very small welfare effects, increasing aggregate welfare by 0.0005%. The reason is that RC renters benefit (0.015%), market renters also capture a minor gain (0.016%), and owners suffer a small loss of welfare (-0.007%). This experiment suggests that changing the spatial allocation of RC housing does not have significant effects on welfare. For our calibration of commuting costs and households' taste for Manhattan, the welfare comparison between the two equilibria reveals the preferences of RC households. They are more concerned with achieving a low rent burden and a larger housing unit than with living in the city center.

Figure 3: Welfare effect of relaxing the main four RC policies: decreasing the mandated fraction of RC sqft by 50% (upper left panel), decreasing the RC discount on market rent by 50% (upper right panel), decreasing the qualifying income cutoff for RC by 50% (lower left panel), and decreasing the maximum size of RC units by 50% (lower right panel). Δ_a are welfare changes measured as permanent consumption changes equivalents, averaged over age a and across owners, renters and RC renters.

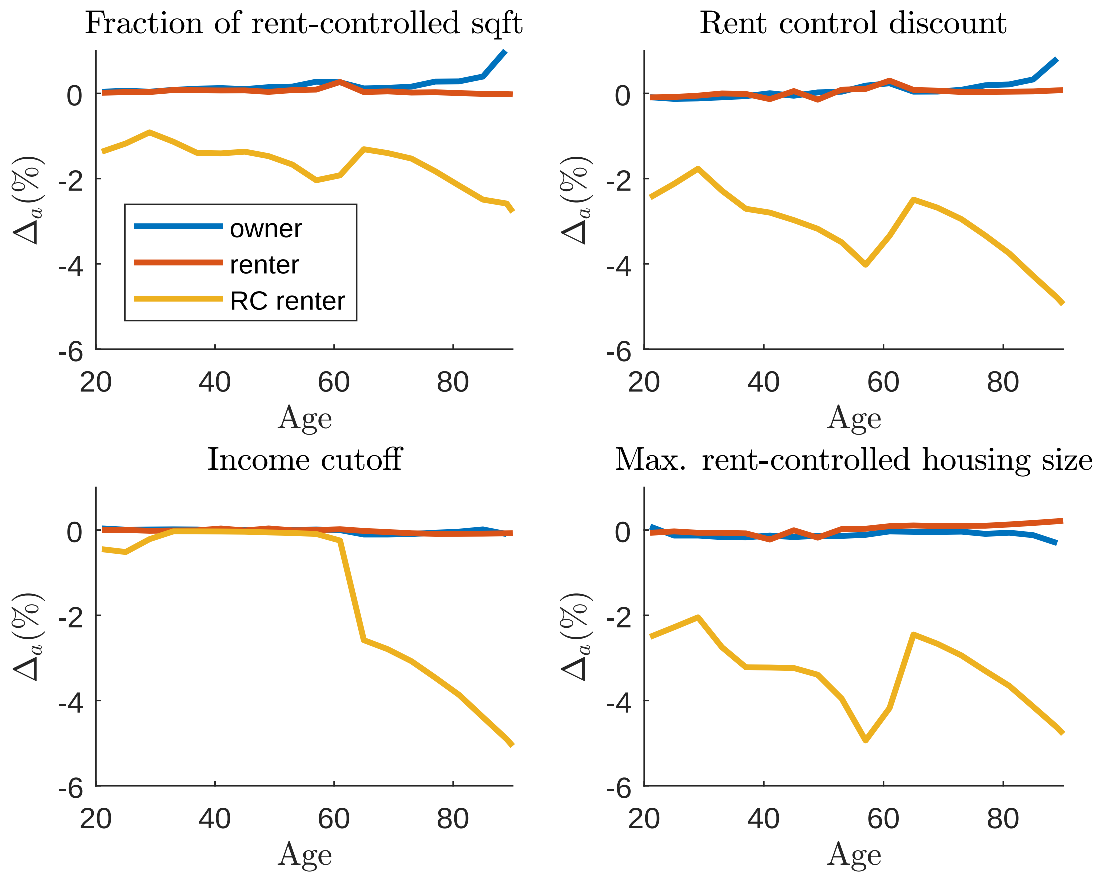

The housing stock available for market renters and owners in Manhattan increases after the policy change. Both alternative uses of space increase. The home ownership rate in Manhattan increases 3%, for example. However, the total housing stock does not rise and because the average housing unit becomes slightly larger, the population share in Manhattan actually falls (-0.4%). Per capita income in zone 1 increases by 0.76%, a combination of higher average productivity (+0.50%) and more hours worked (+0.63%). Rents in zone 1 fall by 0.13%, which combined with the higher income results in more housing affordability in zone 1. The average rent-income ratio falls by 8.6%. The improvement reflects a changing composition of the Manhattan population.

In zone 2, the housing stock falls (-0.14%) as the increase in the share rent control square feet blunts developers' incentives to build. The policy results in a zone-2 housing stock with fewer market rentals (-0.7%) and fewer owner-occupied units (-0.6%). The small reduction in house prices lowers wealth in both zones, and explains why the owners lose from the policy.

5.6 Relaxing Zoning Laws in Zone 1

In the last experiment, we increase $\overline{H}^1$, allowing for more residential housing units in Manhattan. We think of this policy as a relaxation in zoning laws, making it easier to build in Manhattan. The policy leaves the RC parameters unaffected.

Naturally, the housing stock in Manhattan increases. For our parameter choices, it increases by 8.6%. The housing stock in zone 2 falls by 0.16% as developers shift their resources towards zone 1. The overall housing stock in the economy rises by 0.69%. Rents fall by 1.1% in Manhattan and 1.0% in zone 2. Interestingly, the average unit size in zone 1 falls by 23%, while that in zone 2 increases by 4.5%. The extra housing quantity and the smaller unit size allows the population share of Manhattan to rise by 38% from 10.4% of the metro to 14.4%.

Per capita productivity (-31%), hours worked (-9.5%), and income (-35%) in Manhattan all fall substantially with the arrival of new, lower-productivity households. The opposite is true in zone 2. Income inequality within Manhattan rises (+16%) while that within zone 2 falls (-1.8%), as some of the most productive residents of zone 2 move to zone 1, but are still fairly low-productivity compared the the existing residents. This substantially increases average rent to income ratio in zone 1. To that extent, this policy has the same effects as reducing the share of RC units in zone 1.

Lower rents improve housing affordability. The fraction of households in the metro that are rent burdened or severely rent burdened declines by 17% from 19.8 to 16.5%

of the population. Since Manhattan now contains more middle-income households, the fraction of Manhattan households who spend more than 50% of their income on rent actually increases. The opposite is true in zone 2. Relaxing zoning in zone 1 thus appears to have perverse effects, worsening traditional metrics of housing affordability in Manhattan. However, this is due to endogenous migration decisions of households from zone 2 to zone 1 who can now afford Manhattan due to the lower rents. The improvements in housing affordability show up in the zone 2 statistics.

This policy increases the home ownership rate by 1.9% in zone 1 and 4.5% in zone 2. This increase is consistent with the substantial decline in the metro-wide price-income ratio of 7.3%.

Total output falls by 0.09% driven by a decline in non-housing output of 0.18%, and despite a massive increase in construction output in zone 1 (+9%), reflecting the sectoral misallocation of labor. Aggregate hours worked fall by 0.12%. With a lot more households living in Manhattan, commuting time goes down by 5.8%. Households enjoy more leisure time in this economy, causing a first-order increase in their utility. The lower cost of living allows these households to obtain a better "work-life balance", at the expense of a slightly smaller economy.

Figure 4: Welfare effect of reallocating all RC units from zone 1 to zone 2.

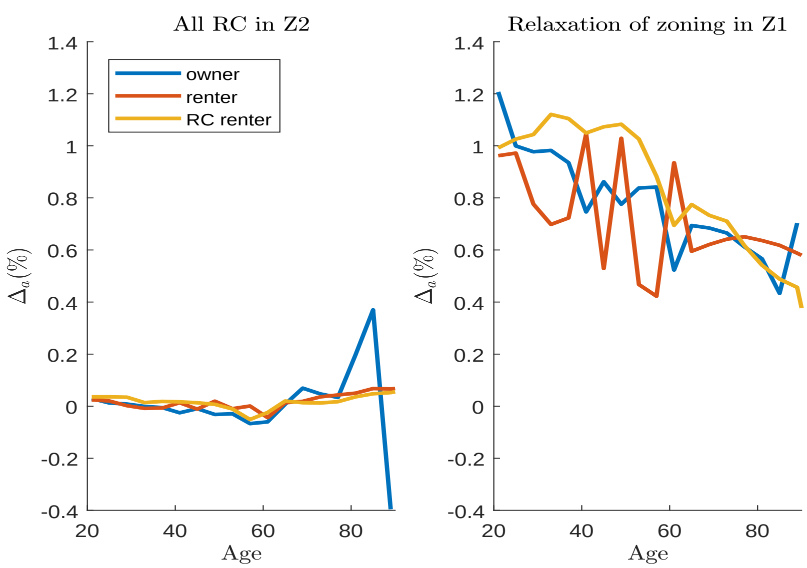

As can be seen in Figure 4, this policy generates the largest welfare gain of 0.85%. It is a Pareto improvement. All groups gain by similar amounts: RC renters (+0.957%), market renters (+0.827%), and owners (+0.855%). The reduction in rents obtained by relaxing zoning also lowers RC rents and generates slightly larger RC units for fewer RC households. Unlike the RC policies, the zoning policy is a progressive policy with gains that are strictly declining in both income and wealth.

6 Conclusion

In a world with rising urbanization rates, the high cost of housing in the major cities of the world has surfaced as a daunting challenge. This paper develops a novel spatial equilibrium model with wealth effects that is amenable to quantitative analysis of housing affordability policies. The model is calibrated to the New York metropolitan area. It matches patterns of wealth accumulation and home ownership over the life-cycle, generates rich income and wealth heterogeneity, and delivers realistic house prices, rents, and wages for Manhattan and the other counties in the metro area. The model features rent control, matching key features of its size and scope in New York, as well as restrictions on residential land use (zoning). We use the model to evaluate various policy changes to the rent control system as well as to zoning, tracing out their aggregate, distributional, and spatial implications.

We find the largest welfare gains for a policy that relaxes zoning laws in Manhattan. A greater supply of housing lowers the cost of living, decreases commuting times and increases leisure. The improved affordability spills over from Manhattan to the rest of the metro area as the population relocates. This policy is a Pareto improvement and progressive in nature. Most changes in rent control policies we evaluate reduce welfare, even as some of them increase the fraction of households in affordable housing units. The best policy for aggregate welfare is one that lowers the fraction of housing that developers must set aside for affordable housing. This policy improves affordability, but hurts renters previously in rent control, and more generally is regressive in nature.

References

Autor, D. H., C. J. Palmer, and P. A. Pathak (2014): "Housing Market Spillovers: Evidence from the End of Rent Control in Cambridge, Massachusetts," *Journal of Political Economy*, 122, 661–717.

——— (2017): "Gentrification and the Amioneity Value of Crime Reductions: Evidence from Rent Deregulation," .

Behrens, K., G. Duranton, and F. Robert-Nicoud (2014): "Productive Cities: Sorting, Selection, and Agglomeration," *Journal of Political Economy*, 122, 507–553.

Berger, D., V. Guerrieri, G. Lorenzoni, and J. Vavra (2015): "House prices and consumer spending," NBER Working Paper Series No. 21667.

Brueckner, J. (1987): *Handbook of Regional and Urban Economics*, Amsterdam: North-Holland, chap. The Structure of Urban Equilibria: A Unified Treatment of the Muth-Mills Model.

Campbell, J. Y. and J. Cocco (2003): "Household Risk Management and Optimal Mortgage Choice," *Quarterly Journal of Economics*, 118, 1449–1494.

Cocco, J. F. (2005): "Portfolio Choice in the Presence of Housing," *Review of Financial Studies*, 18, 535–567.

Davis, M. A. and S. Van Nieuwerburgh (2015): *Handbook of Regional and Urban Economics*, North Holland, chap. Housing, Finance, and the Macroeconomy, Chapter 12.

Diamond, R. and T. McQuade (forthcoming): "Who Wants Aordable Housing in their Backyard? An Equilibrium Analysis of Low Income Property Development," *Journal of Political Economy*.

Diamond, R., T. McQuade, and F. Qian (2017): "The Effects of Rent Control Expansion on Tenants, Landlords, and Inequality: Evidence from San Francisco," .

Eeckhout, J., R. Pinheiro, and K. Schmidheiny (2014): "Spatial Sorting," *Journal of Political Economy*, 122, 554–620.

Favilukis, J., S. C. Ludvigson, and S. Van Nieuwerburgh (2017): "The Macroeconomic Effects of Housing Wealth, Housing Finance and Limited Risk Sharing in General Equilibrium," *Journal of Political Economy*, 125, 140–222.

Gomes, F. and A. Michaelides (2008): "Asset Pricing with Limited Risk Sharing and Heterogenous Agents," *Review of Financial Studies*, 21 (1), 415–449.

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

Heathcote, J., K. Storesletten, and G. L. Violante (2017): "Optimal Tax Progressivity: An Analytical Framework," *Quarterly Journal of Economics*, 132, 1693–1754.

Hizmo, A. (2015): "Risk in Housing Markets: An Equilibrium Approach," Working Paper New York University.

Imrohoroglu, A., K. Matoba, and S. Tuzel (2016): "Proposition 13: An Equilibrium Analysis," Working Paper University of Southern California.

Kaplan, G., K. Mitman, and G. Violante (2016): "Consumption and House Prices in the Great Recession: Model meets Evidence," Working Paper New York University.

Krusell, P. and A. A. Smith (1998): "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy*, 106, 867–896.

Ortalo-Magné, F. and A. Prat (2016): "Spatial Asset Pricing: A First Step," *Economica*, 83, 130–171.

Rappaport, J. (2014): "Monocentric City Redux," Working Paper Federal reserve Bank of Kansas City.

Roback, J. (1982): "Wages, rents, and the quality of life," *The Journal of Political Economy*, 90, 1257.

Rosen, S. (1979): "Wage-based indexes of urban quality of life," *Current issues in urban economics*, 3.

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

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

Sieg, H. and C. Yoon (2017): "Waiting for Affordable Housing," .

Storesletten, K., C. Telmer, and A. Yaron (2006): "Asset Pricing with Idiosyncratic Risk and Overlapping Generations," *Review of Economic Dynamics*, Forthcoming, working Paper, University of Oslo.

Tabellini, G. (1991): "The Politics of Intergenerational Redistribution," *Journal of Political Economy*, 99, 335–357.

Van Nieuwerburgh, S. and P.-O. Weill (2010): "Why Has House Price Dispersion Gone Up?" *Review of Economic Studies*, 77, 1567–1606.

Yao, R. and H. H. Zhang (2004): "Optimal Consumption and Portfolio Choices with Risky Housing and Borrowing Constraint," *Review of Financial Studies*.

A Appendix

A.1 Eliminating idiosyncratic productivity from state space

Suppose that the idiosyncratic productivity process follows a random walk: $z_{t+1} = z_t A_{t+1}$ where A_{t+1} is a random variable which is independent of z_t (we can allow the mean or the variance of A to be agent specific, or to depend on age). Then we will conjecture and verify that the value function is linear in $z_t^{(1-\gamma)(1-\alpha_n)}$: $V(x_t, z_t) = z_t^{(1-\gamma)(1-\alpha_n)} v(xs_t)$ where $xs_t = x_t/z_t$ is the net worth scaled by productivity and where $v(xs_t) = V(x_t, 1)$. We can then rewrite the optimization problem in a way that avoids using z_t. We write the problems with a proportional labor income tax, results carry through with the progressive tax used in the main text.

Suppose that our conjecture is true at $t+1$: $V(x_{t+1}, z_{t+1}, a+1, S_{t+1}) = z_{t+1}^{(1-\gamma)(1-\alpha_n)} v(xs_{t+1}, a+1, S_{t+1})$. Note that this conjecture is true in the last period of a household's life if the household has no bequest motive, or if the bequest is over wealth with the curvature $(1-\gamma)(1-\alpha_n)$.[33] Define $cs_t = c_t/z_t$, $hs_t = h_t/z_t$, $\widehat{hs}_t = \widehat{h}_t/z_t$, and $bs_{t+1} = b_{t+1}/z_t$. Then, it should be straight forward to check that the renter's problem can be rewritten as:

$$\begin{aligned}
&v_{R,\ell}(xs_t, a, S_t) = \max_{cs_t, hs_t, n_t} U(cs_t, hs_t, n_t, \ell, a) + \beta E_t[A_{t+1}^{(1-\gamma)(1-\alpha_n)} v(xs_{t+1}, a+1, S_{t+1})] \text{ s.t.} \\
&cs_t + R_t^\ell hs_t + Q * bs_{t+1} = n_t G^a W_t (1-\tau^{SS}) + xs_t \\
&xs_{t+1} = bs_{t+1}/A_{t+1} \geq 0 \\
&n_t \geq 0
\end{aligned} \tag{10}$$

and the owner's as:

$$\begin{aligned}
&v_{O,\ell}(xs_t, a, S_t) = \max_{cs_t, hs_t, \widehat{hs}_t, n_t} U(cs_t, hs_t, n_t, \ell, a) + \beta E_t[A_{t+1}^{(1-\gamma)(1-\alpha_n)} v(xs_{t+1}, a+1, S_{t+1})] \text{ s.t.} \\
&cs_t + P_t^\ell hs_t + Q * bs_{t+1} + (P_t^\ell - R_t^\ell)\widehat{hs}_t = n_t G^a W_t (1-\tau^{SS}) + xs_t \\
&xs_{t+1} = \left(bs_{t+1} + P_{t+1}^\ell (hs_t + \widehat{hs}_t)(1-\delta-\tau^P) - P_{t+1}^\ell f(\widehat{hs}_t)\right)/A_{t+1} \\
&Q * bs_{t+1} \geq -P_t^\ell(\theta_{res} hs_t + \theta_{inv}\widehat{hs}_t) \\
&n_t \geq 0
\end{aligned} \tag{11}$$

This rescaling is why we require α_n to be constant. If z_t is a stationary process, then we cannot do such a rescaling, and we would not require a constant α_n. Even if z_t is non-stationary, this problem can be solved without rescaling. This is because households have a finite lifespan, therefore z_t is bounded and can be discretized for the numerical solution. However, in both cases, we would need to keep track of z_t as a state variable. Assuming a non-stationary z_t and then rescaling greatly speeds up the numerical procedure.

A.2 Analytical solution for housing and labor supply choices

We will consider the scaled problem, although the same applies to the original problem. We will solve only the worker's problem here. A retiree's problem is analogous, but simpler because there is one fewer choice as $n_t = 0$. For the retirees, out of the four cases described below, only cases one and two are relevant. If z_t is a stationary process, then the scaling is unnecessary; everything in this section applies, but in the choice for hours, G^a must be substituted by $z_t G^a$.

[33] In the first case $V = v = 0$. In the second case $V(x,z) = x^{(1-\gamma)(1-\alpha_n)} = z^{(1-\gamma)(1-\alpha_n)} v(xs)$

First, consider the renter's problem and let λ_t be the Lagrange multiplier on the budget constraint, ν_t be the Lagrange multiplier on the borrowing constraint, and ξ_t be the Lagrange multiplier on the labor constraint. The numerical strategy is to choose cs_t in order to maximize the household's utility. Here we will show that the other choices (n_t and h_t) can be written as analytic functions of cs_t.

Case 1: $\nu_t = 0$ *and* $\xi_t = 0$. In this case the household is unconstrained. The first order conditions are:

$$\begin{aligned} &(1-\gamma)\alpha_{c,a} cs_t^{-1} U_t = \lambda_t \\ &(1-\gamma)\alpha_n (1-n_t-\phi_\ell)^{-1} U_t = \lambda_t G^a W_t (1-\tau^{SS}) \\ &(1-\gamma)\alpha_{h,a} hs_t^{-1} U_t = \lambda_t R_t^\ell \\ &\lambda_t = Q\beta E_t [A_{t+1}^{(1-\gamma)(1-\alpha_n)-1} \frac{\partial v_{t+1}}{\partial xs_{t+1}}] \end{aligned} \tag{12}$$

By rearranging, it is clear that conditional on choosing a location ℓ, $hs_t = \frac{\alpha_{h,a}}{\alpha_{c,a}} \frac{1}{R_t^\ell} cs_t$ and $n_t = 1 - \phi_\ell - \frac{\alpha_n}{\alpha_{c,a}} \frac{1}{G^a W_t (1-\tau^{SS})} cs_t$.

Case 2: $\nu_t > 0$ *and* $\xi_t = 0$. In this case the borrowing constraint binds and $bs_{t+1} = 0$ but the labor constraint does not. The first order conditions in the first three lines of equation 12 are still correct. It is still the case that conditional on choosing a location ℓ, $hs_t = \frac{\alpha_{h,a}}{\alpha_{c,a}} \frac{1}{R_t^\ell} cs_t$ and $n_t = 1 - \phi_\ell - \frac{\alpha_n}{\alpha_{c,a}} \frac{1}{G^a W_t (1-\tau^{SS})} cs_t$. By plugging these into the budget constraint, we can explicitly solve for $cs_t = \alpha_{c,a} \left((1-\phi_\ell) G^a W_t (1-\tau^{SS}) + xs_t \right)$.

Case 3: $\nu_t = 0$ *and* $\xi_t > 0$. In this case the borrowing constraint does not bind, but the labor constraint does, implying $n_t = 0$. The first order conditions in the first, third, and fourth lines of equation 12 are still correct. As in case 1, conditional on choosing a location ℓ, $hs_t = \frac{\alpha_{h,a}}{\alpha_{c,a}} \frac{1}{R_t^\ell} cs_t$.

Case 4: $\nu_t > 0$ *and* $\xi_t > 0$. In this case both constraints bind, implying $n_t = 0$ and $bs_{t+1} = 0$. The first order conditions in the first and third lines of equation 12 are still correct. Now, conditional on choosing a location ℓ, $hs_t = \frac{\alpha_{h,a}}{\alpha_{c,a}} \frac{1}{R_t^\ell} cs_t$. By plugging this into the budget constraint, we can explicitly solve for $cs_t = \frac{\alpha_{c,a}}{\alpha_{c,a}+\alpha_{h,a}} xs_t$.

Next, consider the owner's problem and let λ_t be the Lagrange multiplier on the budget constraint, ν_t be the Lagrange multiplier on the borrowing constraint, and ξ_t be the Lagrange multiplier on the labor constraint. The numerical strategy is to choose cs_t and $\widehat{hs}_t$ in order to maximize the household's utility. Here we will show that the other choices (n_t and h_t) can be written as analytic functions of cs_t and $\widehat{hs}_t$.

Case 1: $\nu_t = 0$ *and* $\xi_t = 0$. In this case the household is unconstrained. The first order conditions are:

$$\begin{aligned} &(1-\gamma)\alpha_{c,a} cs_t^{-1} U_t = \lambda_t \\ &(1-\gamma)\alpha_n (1-n_t-\phi_\ell)^{-1} U_t = \lambda_t G^a W_t (1-\tau^{SS}) \\ &(1-\gamma)\alpha_{h,a} hs_t^{-1} U_t + \beta E_t [A_{t+1}^{(1-\gamma)(1-\alpha_n)-1} \frac{\partial v_{t+1}}{\partial xs_{t+1}} P_{t+1}^\ell (1-\delta-\tau^P)] = \lambda_t P_t^\ell \\ &\beta E_t [A_{t+1}^{(1-\gamma)(1-\alpha_n)-1} \frac{\partial v_{t+1}}{\partial xs_{t+1}} P_{t+1}^\ell (1-\delta-\tau^P - f'(\widehat{h}_t))] = \lambda_t (P_t^\ell - R_t^\ell) \\ &\lambda_t = Q\beta E_t [A_{t+1}^{(1-\gamma)(1-\alpha_n)-1} \frac{\partial v_{t+1}}{\partial xs_{t+1}}] \end{aligned} \tag{13}$$

By rearranging, it is clear that conditional on choosing zone ℓ, $hs_t = \frac{\alpha_{h,a}}{\alpha_{c,a}} \frac{1}{R_t^\ell} cs_t * \frac{1-\delta-\tau^P-f'(\widehat{h}_t)}{1-\delta-\tau^P-f'(\widehat{h}_t)\frac{P_t^\ell}{R_t^\ell}}$ and $n_t = 1 - \phi_\ell - \frac{\alpha_n}{\alpha_{c,a}} \frac{1}{G^a W_t (1-\tau^{SS})} cs_t$.

Case 2: $\nu_t > 0$ *and* $\xi_t = 0$. In this case the borrowing constraint binds implying $bs_{t+1} =$

$-P_t^\ell(\theta_{res}hs_t+\theta_{inv}\widehat{hs}_t)/Q$, but the labor constraint does not bind. Eliminating bs_{t+1} from the budget constraint, we can rewrite it as

$$cs_t + P_t^\ell(1-\theta_{res})hs_t + (P_t^\ell(1-\theta_{inv}) - R_t^\ell)\widehat{hs}_t = n_t G^a W_t(1-\tau^{SS}) + xs_t$$

and resolving gives the following set of first order conditions:

$$\begin{aligned}
&(1-\gamma)\alpha_{c,a}cs_t^{-1}U_t = \lambda_t\\
&(1-\gamma)\alpha_n(1-n_t-\phi_\ell)^{-1}U_t = \lambda_t G^a W_t(1-\tau^{SS})\\
&(1-\gamma)\alpha_{h,a}hs_t^{-1}U_t + \beta E_t[A_{t+1}^{(1-\gamma)(1-\alpha_n)-1}\tfrac{\partial v_{t+1}}{\partial xs_{t+1}}(P_{t+1}^\ell(1-\delta-\tau^P) - P_t^\ell\theta_{res}/Q)] = \lambda_t P_t^\ell(1-\theta_{res})\\
&\beta E_t[A_{t+1}^{(1-\gamma)(1-\alpha_n)-1}\tfrac{\partial v_{t+1}}{\partial xs_{t+1}}(P_{t+1}^\ell(1-\delta-\tau^P-f'(\widehat{h}_t)) - P_t^\ell\theta_{inv}/Q)] = \lambda_t(P_t^\ell * (1-\theta_{inv}) - R_t^\ell)
\end{aligned} \tag{14}$$

The optimal labor choice is the same as in the previous case: $n_t = 1-\phi_\ell - \frac{\alpha_n}{\alpha_{c,a}}\frac{1}{G^a W_t(1-\tau^{SS})}cs_t$. However, the optimal housing choice may now be different. Recall that the numerical strategy is to choose cs_t and $\widehat{hs}_t$. Given those quantities and the binding borrowing constraint, we can use the budget constraint to solve for

$$hs_t = \left(n_t G^a W_t(1-\tau^{SS}) + xs_t - (P_t^\ell(1-\theta_{inv}) - R_t^\ell)\widehat{hs}_t - cs_t\right) / \left(P_t^\ell(1-\theta_{res})\right)$$

Case 3: $\nu_t = 0$ and $\xi_t > 0$. In this case the borrowing constraint does not bind, but the labor constraint does, implying $n_t = 0$. All but the second line of equation 13 are still correct. Conditional on choosing a location ℓ, hs_t is identical to Case 1.

Case 4: $\nu_t > 0$ and $\xi_t > 0$. In this case both constraints bind, implying $n_t = 0$ and $bs_{t+1} = -P_t^\ell(\theta_{res}hs_t+\theta_{inv}\widehat{hs}_t)/Q$. Eliminating bs_{t+1} and n_t from the budget constraint, we can rewrite it as

$$cs_t + P_t^\ell(1-\theta_{res})hs_t + (P_t^\ell(1-\theta_{inv}) - R_t^\ell)\widehat{hs}_t = xs_t$$

We can now solve for hs_t as a function of cs_t and $\widehat{hs}_t$ just as in case 2:

$$hs_t = \left(xs_t - (P_t^\ell(1-\theta_{inv}) - R_t^\ell)\widehat{hs}_t - cs_t\right) / \left(P_t^\ell(1-\theta_{res})\right)$$

A.3 Special case which can be solved analytically

Consider a perpetual renter who is facing a constant wage W and a constant rent R, who is not choosing location, who is not constrained, who faces no idiosyncratic shocks ($A = 1$), and whose productivity and utility are not age dependent ($G^a = 1$, $\alpha_{c,a} = \alpha_c$, and $\alpha_{h,a} = \alpha_h$ $\forall a$). His problem can be written as:

$$\begin{aligned}
&v(xs_t, a) = \max_{cs,hs,n} \tfrac{1}{1-\gamma}\left(cs_t^{\alpha_c}hs_t^{\alpha_h}(1-n_t)^{\alpha_n}\right)^{1-\gamma} + \beta E_t[v(xs_{t+1}, a+1)] \text{ s.t.}\\
&xs_{t+1} = \tfrac{1}{Q}(xs_t + n_t W - cs_t - hs_t R)
\end{aligned} \tag{15}$$

As shown earlier, the optimal housing and labor choices satisfy: $hs_t = \frac{\alpha_h}{\alpha_c}\frac{1}{R}cs_t$ and $n_t = 1 - \frac{\alpha_n}{\alpha_c}\frac{1}{W}cs_t$. Redefining $\widehat{cs} = \frac{1}{\alpha_c}cs$ and plugging these into the maximization problem, the problem is rewritten as:

$$\begin{aligned}
&v(xs_t, a) = \max_{\widehat{cs}} \tfrac{\overline{U}}{1-\gamma}\widehat{cs}_t^{1-\gamma} + \beta E_t[v(xs_{t+1}, a+1)] \text{ s.t.}\\
&xs_{t+1} = \tfrac{1}{Q}(xs_t + W - \widehat{cs}_t)
\end{aligned} \tag{16}$$

where $\overline{U} = \left(\alpha_c^{\alpha_c}\alpha_h^{\alpha_h}\alpha_n^{\alpha_n}R^{-\alpha_h}W^{-\alpha_n}\right)^{1-\gamma}$. Next we can guess and verify that the value function has the form $v(xs_t, a) = \frac{v_a}{1-\gamma} * \left(xs_t + \frac{1}{1-Q_a}W\right)^{1-\gamma}$ where v_a and Q_a are constants that depend on age a. Suppose this is true for $a+1$. Then the problem is:

$$\begin{aligned} v(xs_t, a) &= \max_{\widehat{cs}} \frac{\overline{U}}{1-\gamma}\widehat{cs}_t^{1-\gamma} + \frac{v_{a+1}}{1-\gamma}\beta Q^{-(1-\gamma)}(xs_t + W - \widehat{cs}_t + W\frac{Q}{1-Q_{a+1}})^{1-\gamma} \\ &= \max_{\widehat{cs}} \frac{\overline{U}}{1-\gamma}\widehat{cs}_t^{1-\gamma} + \frac{v_{a+1}}{1-\gamma}\beta Q^{-(1-\gamma)}(xs_t - \widehat{cs}_t + W\frac{1-Q_{a+1}+Q}{1-Q_{a+1}})^{1-\gamma} \end{aligned} \tag{17}$$

Define $X_{a+1} = v_{a+1}Q^{-(1-\gamma)}\beta$. Then the first order condition is: $\overline{U} * \widehat{cs}_t^{-\gamma} = X_{a+1} * (xs_t - \widehat{cs}_t + W\frac{1-Q_{a+1}+Q}{1-Q_{a+1}})^{-\gamma}$. Rearranging, we can solve for optimal consumption:

$$\begin{aligned} \widehat{cs}_t &= \frac{\left(\frac{X_{a+1}}{\overline{U}}\right)^{-1/\gamma}}{1+\left(\frac{X_{a+1}}{\overline{U}}\right)^{-1/\gamma}}\left(xs_t + W\frac{1-Q_{a+1}+Q}{1-Q_{a+1}}\right) \\ xs_{t+1} + \frac{1}{1-Q_{a+1}}W &= \frac{1}{1+\left(\frac{X_{a+1}}{\overline{U}}\right)^{-1/\gamma}}\left(xs_t + W\frac{1-Q_{a+1}+Q}{1-Q_{a+1}}\right) \end{aligned} \tag{18}$$

Plugging this back into the original problem:

$$\begin{aligned} v(xs_t, a) &= \left(\overline{U}\left(\frac{\left(\frac{X_{a+1}}{\overline{U}}\right)^{-1/\gamma}}{1+\left(\frac{X_{a+1}}{\overline{U}}\right)^{-1/\gamma}}\right)^{1-\gamma} + X_{a+1}\left(\frac{1}{1+\left(\frac{X_{a+1}}{\overline{U}}\right)^{-1/\gamma}}\right)^{1-\gamma}\right)\frac{\left(xs_t+W\frac{1-Q_{a+1}+Q}{1-Q_{a+1}}\right)^{1-\gamma}}{1-\gamma} \\ &= \overline{U}\left(1+\left(\frac{X_{a+1}}{\overline{U}}\right)^{-1/\gamma}\right)^{-(1-\gamma)}\left(\left(\frac{X_{a+1}}{\overline{U}}\right)^{-(1-\gamma)/\gamma} + \left(\frac{X_{a+1}}{\overline{U}}\right)\right)\frac{\left(xs_t+W\frac{1-Q_{a+1}+Q}{1-Q_{a+1}}\right)^{1-\gamma}}{1-\gamma} \\ &= X_{a+1}\left(1+\left(\frac{X_{a+1}}{\overline{U}}\right)^{-1/\gamma}\right)^{\gamma}\frac{\left(xs_t+W\frac{1-Q_{a+1}+Q}{1-Q_{a+1}}\right)^{1-\gamma}}{1-\gamma} \end{aligned} \tag{19}$$

This verifies the conjecture. The age dependent constants take the following form:

$$\begin{aligned} v_a &= X_{a+1}\left(1+\left(\frac{X_{a+1}}{\overline{U}}\right)^{-1/\gamma}\right)^{\gamma} \\ &= \beta Q^{-(1-\gamma)}v_{a+1}\left(1+\left(v_{a+1}\beta Q^{-(1-\gamma)}\overline{U}^{-1}\right)^{-1/\gamma}\right)^{\gamma} \\ Q_a &= \frac{Q}{1+Q-Q_{a+1}} \end{aligned} \tag{20}$$

Note that $Q_\infty = Q$ and $v_\infty = \overline{U}\left(1-\beta^{\frac{1}{\gamma}}Q^{\frac{-(1-\gamma)}{\gamma}}\right)^{-\gamma}$.

A.4 Commuting costs and composition of Zone 1

From the household's FOC, we know that $\frac{\partial U}{\partial C} = \frac{\partial U}{\partial N} \times \frac{1}{w}$ where C is the numeraire, N is hours worked, and w is the wage. Suppose that moving one unit of distance towards center decreases the hourly commuting cost by ϕ_T and the financial commuting cost by ϕ_F. Also, suppose that the price is a function of distance from center $P(x)$.

First, consider time costs only ($\phi_F = 0$). The cost of decreasing the commute by d is $d \times H \times P'(x) \times \frac{\partial U}{\partial C}$, this is the amount of housing consumed H, multiplied by the price increase at the current location $P'(x) \times d$, multiplied by the marginal utility of the numeraire good. The benefit of decreasing the commute by d is $d \times \phi_T \times \frac{\partial U}{\partial N} = d \times \phi_T \times w \times \frac{\partial U}{\partial C}$, this is the marginal utility of leisure, multiplied by the extra leisure $d \times \phi_T$. Equating the cost to the benefit and rearranging: $P'(x) = \phi_T \frac{w}{H}$. The left hand side represents one's willingness to pay per square foot implying that agents with high $\frac{w}{H}$ are willing to pay a higher price. For a fixed amount of wealth, high income agents have higher $\frac{w}{H}$ because individual productivity is stationary, therefore high income agents tend to save relatively more and consume relatively less of their wealth ($\frac{w}{H}$ would be constant if individual productivity had permanent shocks). For a fixed income, high wealth agents have higher $\frac{w}{H}$ because, consistent with the Permanent Income Hypothesis, for a fixed w, high wealth agents are willing to spend more on housing.

Next, consider financial costs only ($\phi_T = 0$). The cost of decreasing the commute is the same as before $d \times H \times P'(x) \times \frac{\partial U}{\partial C}$. The benefit of decreasing the commute is $d \times \phi_F \times \frac{\partial U}{\partial C}$, this is the financial saving $d \times \phi_F$ multiplied by the marginal utility of the numeraire. Equating the cost to the benefit: $P'(x) = \phi_F \frac{1}{H}$. Low H agents are willing to pay a higher price. Agents who have low wealth or low income tend to have lower housing demand H and are willing to pay more per square foot to reduce their commute. The intuition is that the financial cost is fixed, thus agents with low housing demand are willing to pay a much higher price per square foot to 'ammortize' the benefit of not paying the fixed cost.

A.5 One-period case which can be solved analytically

There are m agents, m^c consumption producing firms, m^1 construction firms in zone 1, and m^2 construction firms in zone 2. There are two zones with sizes $m\overline{h^1}$ and $m\overline{h}^2$. Agents have initial wealth $W = 0$ and earn a wage w. They live for one period only, and there is no resale value for the housing that they buy.

Conditional on a zone, a household maximizes $U = c^{\alpha_c} h^{\alpha_h} (1 - \lambda - x)^{\alpha_n}$ subject to $c + P * h = W + w * x$ where λ is a zone specific time cost and P is a zone specific housing price ($\lambda = 0$ in zone 1). This can be rewritten as:

$$U = \max_{h,x} (W + w * x - P * h)^{\alpha_c} h^{\alpha_h} (1 - \lambda - x)^{\alpha_n} \tag{21}$$

The first order conditions imply the following solution:

$$\begin{aligned} c &= \alpha_c((1-\lambda)w + W) \\ h &= \alpha_h((1-\lambda)w + W) \\ x &= (\alpha_c + \alpha_h)(1-\lambda) - \alpha_n \frac{W}{w} \\ U &= \left(\frac{1}{P}\right)^{\alpha_h} \left(\frac{1}{w}\right)^{\alpha_n} \alpha_c^{\alpha_c} \alpha_h^{\alpha_h} \alpha_n^{\alpha_n} ((1-\lambda)w + W) \end{aligned} \tag{22}$$

Here we used $\alpha_c + \alpha_h + \alpha_n = 1$.

Each consumption producing firm chooses hours x_c to maximize $\pi_c = x_c^{\rho_c} - w x_c$ which implies that $w = \rho_c x_c^{\rho_c - 1}$. Each construction firm in zone 1 maximizes $\pi_1 = \left(1 - \frac{H^1}{m\overline{h^1}}\right) P_1 x_1^{\rho_h} - w x_1$ which implies that $w = \left(1 - \frac{H^1}{m\overline{h^1}}\right) P_1 \rho_h x_1^{\rho_h - 1}$. Each construction firm in zone 2 maximizes $\pi_2 = \left(1 - \frac{H^2}{m\overline{h^2}}\right) P_2 x_2^{\rho_h} - w x_2$ which implies that $w = \left(1 - \frac{H^2}{m\overline{h^2}}\right) P_2 \rho_h x_2^{\rho_h - 1}$. Here H^1 and H^2 are the total amount of housing built in each zone.

Equilibrium implies that the following equations must be satisfied.

$$P_2 = P_1 (1 - \lambda)^{1/\alpha_h} \tag{23}$$

Equation 23 says that for households to be indifferent between the two zones, their utility of living in each zone must be the same.

$$n_1 = \frac{H^1 P^1}{\alpha_h w} \tag{24}$$

$$n_2 = \frac{H^2 P^2}{\alpha_h w(1-\lambda)} \tag{25}$$

$$n_1 + n_2 = m \tag{26}$$

Equations 24 and 25 say that the total number of households in each zone (N_1 and N_2) must equal to the total housing in each zone, divided by the housing size an agent in that zone would demand. The housing size comes from the solution of the agent's problem. Equation 26 says that the sum of agents living in zones 1 and 2 must equal to the total number of agents.

$$w = \rho_c x_c^{\rho_c - 1} \tag{27}$$

$$w = \left(1 - \frac{H^1}{m\overline{\overline{h^1}}}\right) P_1 \rho_h x_1^{\rho_h - 1} \tag{28}$$

$$w = \left(1 - \frac{H^2}{m\overline{\overline{h^2}}}\right) P_2 \rho_h x_2^{\rho_h - 1} \tag{29}$$

Equations 27, 28, and 29 relate each firm's optimal behavior to the wage.

$$H^1 = \left(1 - \frac{H^1}{m\overline{\overline{h^1}}}\right) m_1 x_1^{\rho_h} \tag{30}$$

$$H^2 = \left(1 - \frac{H^2}{m\overline{\overline{h^2}}}\right) m_1 x_2^{\rho_h} \tag{31}$$

Equations 30 and 31 relate each firm's output to the total output of housing in each zone. They can be rewritten as $H^1 = \frac{m\overline{\overline{h^1}} m_1 x_1^{\rho_h}}{m\overline{\overline{h^1}} + m_1 x_1^{\rho_h}}$ and $H^2 = \frac{m\overline{\overline{h^2}} m_2 x_2^{\rho_h}}{m\overline{\overline{h^2}} + m_2 x_2^{\rho_h}}$.

$$(\alpha_c + \alpha_h)(n_1 + n_2(1-\lambda)) = m_c x_c + m_1 x_1 + m_2 x_2 \tag{32}$$

Equation 32 relates labor supply, on the left side, to labor demand, on the right side.

This is 10 equations and 10 unknowns: prices P_1, P_2; labor demand for each firm type x_1, x_2, x_c; number of households living in each zone n_1, n_2; total housing in each zone H^1, H^2; and the wage w. This can can be reduced to a single equation.

First, plug H and P into equations (28) and (29): $w = P_1 \rho_h \frac{m\overline{\overline{h^1}} x_1^{\rho_h - 1}}{m\overline{\overline{h^1}} + m_1 x_1^{\rho_h}} = P_2 \rho_h \frac{m\overline{\overline{h^2}} x_2^{\rho_h - 1}}{m\overline{\overline{h^2}} + m_2 x_2^{\rho_h}}$

Second, plug the wage into equations (24) and (25): $n_1 = \frac{m_1 x_1}{\alpha_h \rho_h}$ and $n_2 = \frac{m_2 x_2}{\alpha_h \rho_h (1-\lambda)}$.

Third, plug n_1 and n_2 into equation (26) to solve for x_2 in terms of x_1: $x_2 = \frac{1-\lambda}{m_2}(m\alpha_h \rho_h - m_1 x_1) = A_0 + A_1 x_1$ where $A_0 = \frac{1-\lambda}{m_2} m \alpha_h \rho_h$ and $A_1 = -m_1 \frac{1-\lambda}{m_2}$.

Fourth, plug $x_2 = A_0 + A_1 x_1$ into the equality between zone 1 and zone 2 firms' wages derived earlier and use equation (23) to get rid of prices: $\frac{m\overline{\overline{h^1}} x_1^{\rho_h - 1}}{m\overline{\overline{h^1}} + m_1 x_1^{\rho_h}} = (1-\lambda)^{1/\alpha_h} \frac{m\overline{\overline{h^2}} (A_0 + A_1 x_1)^{\rho_h - 1}}{m\overline{\overline{h^2}} + m_2 (A_0 + A_1 x_1)^{\rho_h}}$. This is now one equation with one unknown and can be solved numerically.

Fifth, once we have x_1 we can immediately calculate x_2, n_1, n_2, H^1, H^2 but we still need to solve for w and P_1. We can solve for w as a function of P_1 using equation (28). We can then solve for x_c as a function of P_1 using equation (27). We can then plug this into equation (32) to solve for

P_1.

B Data Appendix: New York

B.1 The New York Metro Area

U.S. Office of Management and Budget publishes the list and delineations of Metropolitan Statistical Areas (MSAs) on the Census website (https://www.census.gov/population/metro/data/metrodef.html). The current delineation is as of July 2015. New York-Newark-Jersey City, NY-NJ-PA MSA (NYC MSA) is the most populous MSA among the 382 MSAs in the nation.

NYC MSA consists of 4 metropolitan divisions and 25 counties, spanning three states around New York City. The complete list of counties with state and zone information is presented in Table 4. As previously defined, only New York County (Manhattan borough) is categorized as zone 1 and the rest 24 counties are categorized as zone 2. For informational purposes, the five counties of New York City are appended with parenthesized borough names used in New York City.

Table 4: NYC MSA

County	**State**	**Zone**
New York (Manhattan)	NY	Zone 1
Bergen	NJ	Zone 2
Bronx (Bronx)	NY	Zone 2
Dutchess	NY	Zone 2
Essex	NJ	Zone 2
Hudson	NJ	Zone 2
Hunterdon	NJ	Zone 2
Kings (Brooklyn)	NY	Zone 2
Middlesex	NJ	Zone 2
Monmouth	NJ	Zone 2
Morris	NJ	Zone 2
Nassau	NY	Zone 2
Ocean	NJ	Zone 2
Orange	NY	Zone 2
Passaic	NJ	Zone 2
Pike	PA	Zone 2
Putnam	NY	Zone 2
Queens (Queens)	NY	Zone 2
Richmond (Staten Island)	NY	Zone 2
Rockland	NY	Zone 2
Somerset	NJ	Zone 2
Suffolk	NY	Zone 2
Sussex	NJ	Zone 2
Union	NJ	Zone 2
Westchester	NY	Zone 2

B.2 Population, Housing Stock, and Land Area

The main source for population, housing stock and land area is US Census Bureau `American FactFinder` (`http://factfiner.census.gov`). `American FactFinder` provides comprehensive survey data on a wide range of demographic and housing topics. Using the Advanced Search option on the webpage, topics such as population and housing can be queried alongside geographic filters. We select the DP02 table (selected social characteristics) for population estimates, the DP04 table (selected housing characteristics) for housing estimates, and the GCT-PH1 table (population, housing units, area and density) for land area information. Adding 25 counties separately in the geographic filter, all queried information is retrieved at the county level. We then aggregate the 24 columns as a single zone 2 column.

Since the ACS (American Community Survey) surveys are conducted regularly, the survey year must be additionally specified. We use the 2015 1-year ACS dataset as it contains the most up-to-date numbers available. For Pike County, PA, the 2015 ACS data is not available and we use the 2014 5-year ACS number instead. Given that Pike County accounts only for 0.3% of zone 2 population, the effect of using lagged numbers for Pike County is minimal.

The ratio of the land mass of zone 1 (Manhattan) to the land mass of zone 2 (the other 24 counties of the NY MSA) is 0.0028. However, that ratio is not the appropriate measure of the relative maximum availability of housing in each of the zones since Manhattan zoning allows for taller buildings, smaller lot sizes, etc.

Data on the maximum buildable residential area are graciously computed and shared by Chamna Yoon from Baruch College. He combines the maximum allowed floor area ratio (FAR) to each parcel to construct the maximum residential area for each of the five counties (boroughs) that make up New York City. Manhattan has a maximum residential area of 1,812,692,477 square feet. This is our measure for $\bar{H}^1$. The other four boroughs of NYC combine for a maximum buildable residential area of 4,870,924,726 square feet. Using the land area of each of the boroughs (expressed in square feet), we can calculate the ratio of maximum buildable residential area (sqft) to the land area (sqft). For Manhattan, this number is 2.85. For the other four boroughs of NYC it is 0.62. For Staten Island, the most suburban of the boroughs, it is 0.32. We assume that the Staten Island ratio is representative of the 20 counties in the New York MSA that lie outside NYC since these are more suburban. Applying this ratio to their land area of 222,808,633,344 square feet, this delivers a maximum buildable residential square feet for those 20 counties of 71,305,449,967 square feet. Combining that with the four NYC counties in zone 2, we get a maximum buildable residential area for zone 2 of 76,176,377,693 square feet. This is $\bar{H}^2$. The ratio $\bar{H}^1/\bar{H}^2$ is 0.0238. We argue that this ratio better reflects the relative scarcity of space in Manhattan than the corresponding land mass ratio.

B.3 Income

The main source for the income distribution data is again US Census Bureau `American FactFinder`. From table DP03 (selected economic characteristics), we retrieve the number of households in each of 10 income brackets, ranging from "less than $10,000" for the lowest to "$200,000 or more" for the highest bracket. The distribution suffers from top-coding problem, so we additionally estimate the conditional means for the households in each income bracket. For the eight income brackets except for the lowest and the highest, we simply assume the midpoint of the interval as the conditional mean. For example, for the households in $50,000 to $74,999 bracket, the conditional mean income is assumed to be $62,500. For the lowest bracket, (less than $10,000) we assume the conditional mean is $7,500. Then we can calculate the conditional mean of the highest income bracket, using the average household income and conditional means of the other brackets, since the reported unconditional mean is based on all data.

Our concept of income is household income before taxes. It includes income from (i) wages, salaries, commissions, bonuses, or tips, (ii) farm or non-farm business, proprietorship, or partnership, (iii) social security and railroad retirment payments, (iv) retirement, survivor, or disability pensions, (v) SSI, TANF, family assistance, safety net, other public assistance, or public welfare, (vi) interest, dividends, royalties, estates and trusts.

We aggregate the county-level income distribution into a zone 2 income distribution in two steps. First, the aggregate number of households included in each income bracket is the simple sum of county-level household numbers in the bracket. Second, we calculate the zone 2 conditional mean of the income brackets using the weighted average methods. For the lower nine income brackets, the conditional means are assumed to be constant across counties, so zone 2 conditional means are also the same. For the highest income bracket, we use the county-specific conditional mean of the highest bracket, and calculate its weighted average over the 24 counties. Using these conditional means, and the household distribution over 10 income brackets, the zone 2 average household income can be calculated.

B.4 House Prices, Rental Prices, and Home Ownership

Housing prices and rental prices data come from `Zillow` (http://www.zillow.com/research/data) indices. Zillow publishes Zillow Home Value Index (ZHVI) and Zillow Rent Index (ZRI) monthly. The main advantage of using Zillow indices compared to other indices is that it overcomes sales-composition bias by constantly estimating hypothetical market prices, controlling for hedonics such as house size. We use 2015 year-end data to be consistent with the ACS dataset. There are a few missing counties in ZHVI and ZRI. For the five counties with missing ZHVI index price, we search those counties from `Zillow` (http://www.zillow.com) website, and use the median listing prices instead. For the two counties with missing ZRI index price, we estimate the rents using the price/rent ratio of comparable counties.

Home ownership data is directly from `American FactFinder`. In table DP04 (selected housing characteristics), the *Total housing units* number is divided by *Occupied housing units* and *Vacant housing units*. *Occupied housing units* are further classified into *Owner-occupied* and *Renter-occupied* housing units, which enables us to calculate the home ownership ratio.

B.5 Rent Regulation

The main source for rent regulation data is US Census Bureau New York City Housing and Vacancy Survey (NYCHVS; http://www.census.gov/housing/nychvs). NYCHVS is conducted every three years to comply with New York state and New York City's rent regulation laws. We use the 2014 survey data table, which is the most recent survey data. In Series IA table 14, the number of housing units under various rent-control regulations are available for each of the five NYC boroughs. We define rent-regulated units as those units that are (i) rent controlled, (ii) public housing, (iii) Mitchell Lama housing, (iv) all other government-assisted or regulated housing.

We exclude rent-stabilized units from our definition. Rent stabilized units are restricted in terms of their annual rent increases. The vast majority of units built after 1947 that are rent stabilized are so voluntarily. They receive tax abatements in lieu of subjecting their property to rent stabilization for a defined period of time. Both rent levels and income levels of tenants in rent-stabilized units are in between those of rent-regulated and unregulated units.

We calculate the proportion of rent-regulated units among all the renter-occupied units. The proportion is 16.9% for Manhattan and 13.2% for the other four NYC boroughs.

We use a different data source for the other 20 counties outside of New York City. `Affordable Housing Online` (http://affordablehousingonline.com) provides various rent-related statistics at the county level. For each of the 20 counties outside NYC, we calculate the fraction of rent-regulated units by dividing *Federally Assisted Units* number by *Renter Households* number reported on each county's webpage. We then multiply these %-numbers with the renter-occupied units in ACS data set to calculate the rent-regulated units for the 20 counties. Along with the NYCHVS numbers for the four NYC boroughs, we can aggregate all the 24 counties in zone 2 to calculate the fraction of rent-regulated units. The four NYC boroughs have 1.53 million renter-occupied housing units while the rest of zone 2 has 1.30 million. The resulting fraction of rent-regulated units in zone 2 is 10.4%.

From the NYCHVS, we also calculate the percentage difference in average rent in New York City between our definition of regulated rentals and the others (unregulated plus rent-stabilized). That percentage difference is 49.9%. We apply the same percentage difference to all of the MSA in our model.

Finally, we calculate the percentage difference in average household income (Series IA - Table 9) in New York City between our definition of regulated rentals and the others (unregulated plus rent-stabilized). That percentage difference is 54.2%. This is a moment we can compute in the model and compare to the data.

B.6 Migration

We use county-to-county migration data for 2006-2010 and 2010-2014 from the 5-year American Community Survey for the 25 counties in the New York metropolitan area. For each county and survey wave, we compute net migration rates (inflow minus outflow divided by population). When one person enters the New York labor market and another one leaves, the model is unchanged, so net migration is the relevant concept for the model. We aggregate net migration for the 24 counties in zone 2. The net migration rate over the 5-year period between 2010-2014 for the entire MSA is -0.15%, or -0.03% per year. First, this net migration rate is minuscule: only about 30,000 people moved in over a 5-year period on a MSA population of 20 million. Of course, this masks much larger gross flows: about 824,000 came into the MSA and 854,000 left. Second, Manhattan (zone 1) saw a net inflow of 30,000 people coming from outside the MSA while the rest of the MSA (zone 2) saw a net outflow to the rest of the country/world of 60,000. This is the opposite pattern than what we would expect if the OOT purchases prompted migration of residents, since OOT purchases were much stronger in Manhattan than in the rest of the MSA (twice as large). Third, comparing the net migration in the 2010-2014 period to that in the 2006-2010 period, we find that the net migration rate rose, from -73,000 to -30,000. The net migration rate rises from -0.38% in 2006-2010 to -0.15% in the 2010-2014 period. The rise in OOT purchases over time did not coincide with a decline in net migration, but with an increase. In other words, not only are the relevant net migration rates tiny, they also have the wrong-sign cross-sectional correlation with the spatial OOT pattern, and with the time-series of OOT purchases. We conclude that there is little evidence in the New York data of substantial net migration responses to OOT purchases.

C Labor Income Calibration

Before-tax labor income for household i of age a is given by:

$$y_{i,a} = W_t n_t^i G^a z^i$$

where G^a is a function of age and z^i is the idiosyncratic component of productivity. For purposes of calibrating the productivity process, we assume that labor supply is inelastic.

We set the productivity states are $z^i \in Z = [0.255, 0.753, 1.453, 3.522]$ to match the observed mean NY household income, scaled by the NY metro area average, in the income groups below \$50,000, between \$50,000 and \$100,000, between \$100,000 and \$200,000, and above \$200,000. The NY income data is top-coded. For each county in the NY metro area, we observe the number of households whose earnings exceed \$200,000. Because we also observe average earnings (without top-coding), we can infer the average income of those in the top coded group.

The transition probability matrix for z is $\mathbb{P}$ for β^L agents. We impose the following restrictions:

$$\mathbb{P} = \begin{bmatrix} p_{11} & 1-p_{11} & 0 & 0 \\ (1-p_{22})/2 & p_{22} & (1-p_{22})/2 & 0 \\ 0 & (1-p_{33})/2 & p_{33} & (1-p_{33})/2 \\ 0 & 0 & 1-p_{44} & p_{44} \end{bmatrix}$$

For β^H types, the transition probability matrix is the same, except for the last two entries which are $1 - p_{44} - p^H$ and $p_{44} + p^H$, where $p^H < 1 - p_{44}$. We pin down the five parameters

$$(p_{11}, p_{22}, p_{33}, p_{44}, p^H) = (.933, .806, .780, .632, .238)$$

to match the following five moments. We match the population shares in each of the four income groups defined above: 39.1%, 27.2%, 23.7%, and 10.0%, respectively. Given that population shares sum to one, that delivers three moments. We match the persistence of individual labor income to a value of 0.9, based on evidence form the PSID in Storesletten, Telmer, and Yaron (2006). Finally, we choose p^H to match the fraction of high-wealth households in the top 10% of the income distribution.

We choose $G^a = G_0^a G_1^a$ to capture both how average labor income changes with age and how the dispersion of income changes with age. The first component G_0^a captures the mean shift, while G_1^a captures the dispersion. We use data from the SCF to estimate the function G^a; we use all ten waves from 1983 to 2010. The procedure is as follows. First, we pool all households in the SCF into four-year age bins; let there be A such bins . We calculate average real income by age group using SCF weights. We average across the ten SCF waves. Then, we estimate a regression of log average labor income by age on a constant and log age. Denote the exponential of the fitted value by Y_a, the average income by age. We define $G_0 = Y_a / \mathbb{E}_a[Y_a]$ to be deterministic life cycle profile for mean earnings, normalized to be mean 1 over the life-cycle; $\mathbb{E}_a[Y_a]$ is the average using the age shares in the population. G_0^a is a $A \times 1$ vector.

Second, in each SCF age group, we further sort households into four groups based on the pre-tax real labor income. The income groups are the same as in the NYC data: below \$50,000, between \$50,000 and \$100,000, between \$100,000 and \$200,000, and above \$200,000. The bin cutoffs are adjusted for inflation for the earlier SCF waves. We form average income conditional on age and income, again averaging over survey waves. We then estimate a regression of log average income by age on a constant and log age, separately for each of the four income groups. Denote the exponential of the fitted value by Y_{ia}, the average income by age conditional on being in income

Figure 5: Labor Income Calibration

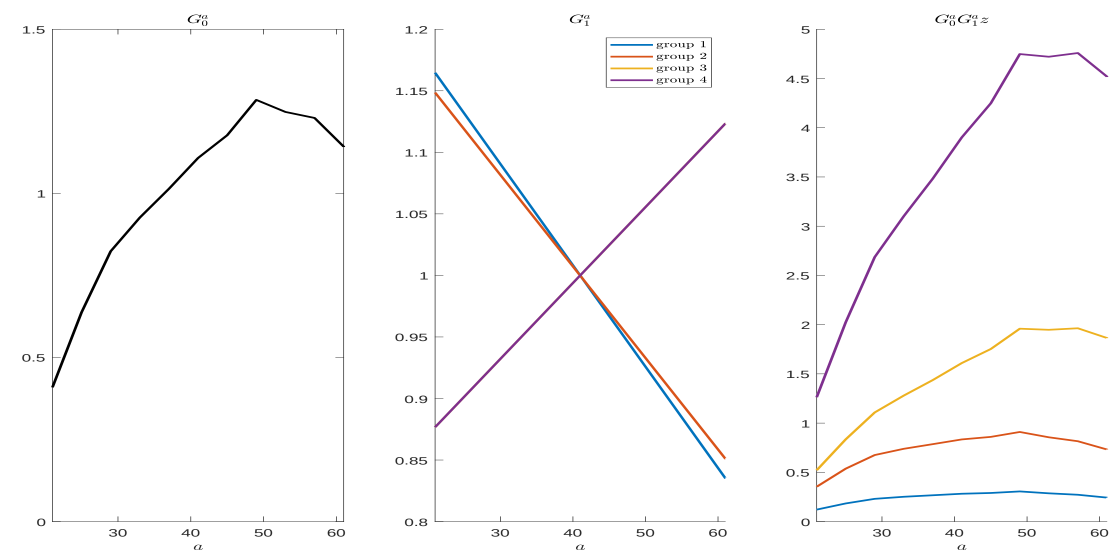

group $i \in \{1, 2, 3, 4\}$. We then construct a $A \times 4$ matrix G_1^a with entries $g_{a,i}$ given by:

$$g_{a,i} = \frac{Y_{ia}}{Y_a z_i c_a}$$

The rows of G_1 describe the income dispersion for a given age a beyond what is already induced by the life-cycle profile for mean earnings Y_a and by the NYC productivity grid z. We scale the four elements of each row (age), such that they average to 1. This is accomplished by the age-specific scalar c_a, where the average takes into account the shares of each income group at that age. The G_1^a adjustment implements a mean-preserving spread of productivity which is age-specific.

The left panel of Figure 5 plots the mean earnings profile G_0^a, which shows the familiar hump shape over the life-cycle; G_a^0 is zero (0.01) for age 65 and later. The middle panel plots the earnings dispersion profile G_1^a. It shows that at young ages, dispersion is dampened since the values for the bottom two groups are above 1 and those for the high income groups below 1. This pattern reverses gradually, implying rising variance of income over the life-cycle. The right panel plots the product, $G_0^a G_1^a z$, which are the efficiency units of labor for the various productivity groups.

Average labor income in the model for working-age households is 0.370. When fitting an AR(1) to labor income, it has an autocorrelation of 0.682 per four years or 0.909 annualized. The innovation standard deviation is 0.197 per four years or 0.099 per year. These are typical values used in the literature. Labor income is higher on average for β^H- than for β^L-types (0.455 vs. 0.342), slightly more persistent (0.687 vs. 0.680) and slightly less volatile (0.196 vs.0.197).

Since we calibrated the productivity process in the model to observed household income in the data, we check how close household income in the model is to the data. This takes into account the endogenous labor supply choice. Every period, we compute total income as pre-tax labor income when the agent is a worker, and as retirement income when she is a retiree. Then we normalize those household income levels by average household income this period. Then, we compute average normalized household income in each of the four percentile ranges (income groups). This gives us a four-point grid every period. Finally, we average those grids over time. This delivers an income grid with values [0.092, 0.530, 1.421, 4.818]. The corresponding grid in the NYC data is [0.255, 0.753, 1.453, 3.522]. This suggests our baseline productivity grid has too much dispersion. Since endogenous labor supply decisions depend on all other parameters and state variables of the model, exactly matching income in model and data is a non-trivial task.

D Housing Supply Elasticity calibration

We compute the long-run housing supply elasticity. It measures what happens to the housing quantity and housing investment in response to a 1% permanent increase in house prices. Define housing investment for a given zone, dropping the location superscript since the treatment is parallel for both zones, as:

$$Y_t^h = \left(1 - \frac{H_{t-1}}{\overline{H}}\right) N_t^{\rho_h}.$$

Note that $H_{t+1} = (1-\delta)H_t + Y_t^h$, so that in steady state, $Y^h = \delta H$. Rewriting the steady state housing investment equation in terms of equilibrium quantities using (5) delivers:

$$H = \frac{1}{\delta}\left(1 - \frac{H}{\overline{\overline{H}}}\right)^{\frac{1}{1-\rho_h}} \rho_h^{\frac{\rho_h}{1-\rho_h}} \overline{P}^{\frac{\rho_h}{1-\rho_h}} W^{\frac{-\rho_h}{1-\rho_h}}$$

Rewrite in logs, using lowercase letters to denote logs:

$$h = -\log(\delta) + \frac{1}{1-\rho_h}\log(1-\exp(h-\overline{h})) + \frac{\rho_h}{1-\rho_h}\overline{p} - \frac{\rho_h}{1-\rho_h}w$$

Rearrange and substitute for $\overline{p}$ in terms of the market price $\overline{p} = \log(ho + (1-ho)\kappa_4) + p$:

$$p = \frac{1-\rho_h}{\rho_h}h - \frac{1}{\rho_h}\log(1-\exp(h-\overline{h})) + k$$

where

$$k = \frac{1-\rho_h}{\rho_h}\log(\delta) + w - \log(ho + (1-ho)\kappa_4)$$

Now take the partial derivative of p w.r.t. h:

$$\frac{\partial p}{\partial h} = \frac{1-\rho_h}{\rho_h} + \frac{1}{\rho_h}\frac{\exp(h-\overline{h})}{1-\exp(h-\overline{h})} + \frac{\partial k}{\partial h}$$

Invert this expression delivers the housing supply elasticity:

$$\frac{\partial h}{\partial p} = \frac{\rho_h}{1-\rho_h + \frac{\exp(h-\overline{h})}{1-\exp(h-\overline{h})} + \left[\rho_h\frac{\partial w}{\partial h} - \rho_h\frac{(1-\kappa_4)}{ho+(1-ho)\kappa_4}\frac{\partial ho}{\partial h}\right]}$$

If (i) the elasticity of wages to housing supply is small ($\frac{\partial w}{\partial h} \approx 0$) and either the rent control distortions are small ($\kappa_4 \approx 0$) or the home ownership rate is inelastic to the housing supply ($\frac{\partial ho}{\partial h} \approx 0$), or (ii) if the two terms in square brackets are positive but approximately cancel each other out, then the last two terms are small. In that case, the housing supply elasticity simplifies to:

$$\frac{\partial h}{\partial p} \approx \frac{\rho_h}{1 - \rho_h + \frac{\exp(h-\overline{h})}{1-\exp(h-\overline{h})}}$$

We will use this approximation to calibrate the housing supply elasticity to the data. Since, in equilibrium, $Y^h = \delta H$, $\partial y^h / \partial p = \partial h / \partial p$.

Note that $h - \overline{h}$ measures how far the housing stock is from the constraint, in percentage terms. As H approaches $\overline{H}$, the term $\frac{\exp(h-\overline{h})}{1-\exp(h-\overline{h})}$ approaches $+\infty$ and the elasticity approaches zero. This is the case in zone 1 for our calibration. If H is far below $\overline{H}$, that term is close to zero and the housing supply elasticity is close to $\frac{\rho_h}{1-\rho_h}$. That is the case for zone 2 in our calibration. Since zone 2 is by far the largest component of the New York metro housing stock, zone 2 dominates the overall housing supply elasticity we calibrate to.