Out-of-town Home Buyers and City Welfare *

Jack Favilukis[†] Stijn Van Nieuwerburgh[‡]

October 2, 2018

Abstract

Many of the world's major cities have attracted a flurry of out-of-town (OOT) home buyers. Such capital inflows affect house prices, rents, construction, labor income, wealth, and ultimately welfare. We develop an equilibrium model, calibrated to the typical U.S. metropolitan area, to quantify the welfare effects of OOT home buyers. When OOT investors buy 10% of the city's housing, welfare among residents falls by 0.74% in consumption equivalent units. Housing becomes less affordable, with rents increasing by 19% and house prices by 10% in our baseline model. A construction boom pushes up city-wide wages, reducing the competitiveness of the city and aggregate employment. The model's ability to generate substantial heterogeneity in income, wealth, and tenure status among residents is crucial for accurately measuring welfare effects of the OOT shock. Policies like taxing OOT buyers to finance local public goods, mandating them to rent out their property, or expanding land available for development can mitigate or reverse welfare losses.

JEL classification: R10, R20, R30, R40, R51, G11, G12, H41, H70, J61
Keywords: house prices, foreign investors, property taxes

*First draft: November 18, 2016. We thank seminar and conference participants at the University of British Columbia Sauder school, NYU Stern CREFR, Imperial College, the American Economic Association in Chicago, the New York City area real estate conference at CUNY Baruch, the AREUEA national conference in Washington D.C, the Bank of England housing conference, the Society for Economic Dynamics conference in Edinburgh, the pre-Western Finance Association real estate conference in Whistler, and the National Bureau of Economic Research Summer Institute REPE. We thank Hae-Kang Lee for excellent research assistance. We thank Kristine Yao and Sam Khater from Core Logic, Chamna Yoon from CUNY Baruch, and Arpit Gupta from NYU Stern for sharing data for the New York metropolitan area and Christophe Spanjaers for sharing data on Paris. We thank Morris Davis, Vadim Elenev, Daniel Greenwald, Arpit Gupta, Adam Guren, Erik Hurst, Yishen Liu, Holger Mueller, Esteban Rossi-Hansberg, Jordan Rappaport, Holger Sieg, Johannes Stroebel, and Andrea Tambalotti for comments.

[†]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 Finance, Graduate School of Business, Columbia University, Uris Hall 809, 3022 Broadway, New York, NY 10027, USA; NBER; CEPR; svnieuwe@gsb.columbia.edu; `https://www0.gsb.columbia.edu/faculty/svannieuwerburgh`.

1 Introduction

Residential investment in major urban centers by out-of-town (OOT), including foreign, investors has been on the rise. For example, the U.S. attracted $153 billion in foreign residential real estate investment between April 2016 and March 2017, accounting for 10% of all purchases. These investment flows are controversial since they tend to leave highly desirable real estate under-utilized and are thought to push up the cost of housing for local residents, fueling the affordability issues these cities already struggle with. Public opposition to OOT investor flows has led many large cities to tax residential property purchases by non-residents. But OOT investors bring benefits as well. They spur new construction which requires local labor. This demand for construction services (and for other non-tradables like restaurants) increases wages for all local workers. They generate additional property tax revenue which can pay for education and infrastructure benefitting local residents. And most importantly, they increase property values for local home owners. The net effect of these costs and benefits is unknown. This paper asks whether OOT buyers of local real estate are good or bad for the city's overall welfare, or equivalently, whether taxes on OOT buyers are welfare improving.

We set up and solve an equilibrium model with the main ingredients necessary to address these questions. 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, and tenure status (own or rent). Home owners also choose how much rental property to own and rent out to local renters. The model generates a cross-sectional distribution over age, labor income, tenure status, housing wealth, and financial wealth. The model delivers realistic wealth accumulation and home ownership patterns over the life-cycle and in the cross-section. The city produces tradable goods and residential housing. While interest rates and goods prices are taken as given, wages, house prices, and rents are determined in the city's equilibrium. This richness is necessary to identify the winners and losers from the OOT investment and to quantify their gains and losses.

We shock this city with an inflow of OOT real estate investors. OOT investors use the home as a pied-à-terre rather than renting it out to locals, thereby removing housing from the market. From the perspective of the locals, the OOT housing demand shock is a source of aggregate risk. Local residents form beliefs over the expected duration of low- and high-OOT investment spells. We calculate how much each local resident would be willing to pay to avoid a transition from the low-OOT state, with zero OOT demand, to the high-OOT housing demand state, which has OOT investors accounting for 10% of housing. This 10% number corresponds to the OOT demand for residential housing in Vancouver in the years leading up to the introduction of the transaction tax in August 2016. Section 2 presents stylized facts on OOT demand suggesting that the 10% share is not uncommon among the major cities in the world.

Given imperfectly elastic housing supply, we find that an inflow of OOT investors is only partially absorbed by new construction. Residential construction increases by 17.3% in the short-run, and eventually increases to accommodate 28% of the OOT demand. Space constraints make housing supply less than perfectly elastic. Rents must rise to induce residents to reduce housing demand so that the market for space can clear. Rents go up by 18.8% in the short-run and by 13.6% in the long-run, implying a substantial increase in the cost of housing. Likewise, house prices increase substantially (+10.0% in short-run, +7.1% in long-run), but by less than rents, because house prices capitalize a possible future reversal to the low OOT state associated with low rents. City-wide wages rise (+1.1%) due to the boom in the construction sector. The aggregate increase in wages reduces the competitiveness of the city and aggregate hours worked fall by 1.1%. Because average earnings are essentially unchanged, the rent-income and price-income ratios rise by the same amount as rents and prices, indicating severe reductions in housing affordability.

Renters, whose cost of living rises, suffer welfare losses from the OOT demand shock. The average renter would be willing to forgo 2.42% of annual consumption to avoid it. Younger renters and older renters, who tend to be poorer, suffer the largest welfare reduction. Owners, on the other hand, benefit not only from the higher wages but also from the capital gains on housing. Younger owners however face higher future housing costs, reducing some of the benefit. The average owner's welfare gain is 0.56%. Aggregating the welfare of all local residents, city welfare falls by 0.74% following a 10% increase in OOT demand.

We investigate how sensitive the aggregate welfare results are to changes in the model parameters, including those that govern the OOT demand. A key assumption is that OOT investors do not rent out their properties to locals. A model where they rent out all of their housing results in welfare effects that are approximately zero. Differences in the price elasticity of demand between OOT and local property investors are quantitatively not important. The welfare effects are driven by the effective removal of a part of the housing stock from local residents' use. A policy that encourages or mandates OOT owners to rent out their property at least part of the year would be welfare increasing. Welfare costs of OOT demand are nearly proportional to the size of the OOT demand shock, and only modestly lower when the OOT shock is less persistent. The welfare costs are larger when housing supply elasticity is lower, suggesting that the cost in cities like New York, San Francisco, or Miami may be substantially higher. Costs can also be substantially higher when the housing demand elasticity is lower, a parameter for which the literature has considered a wide range of values. For low but still plausible values of the housing demand elasticity, rent increases can be as large as 30% and price increases as large as 15%. We find that the welfare cost is substantially underestimated when the model generates too little heterogeneity in wealth. Related, the welfare cost falls substantially when there is no minimum housing size, (counter-factually) allowing renters to live in very small quarters. Policies that reduce minimum size restrictions may mitigate the welfare cost of OOT investors. In sum, all model variants imply robustly negative welfare costs. Only an economy where all residents are forced to be owners and that understates wealth inequality can generate a (small) aggregate welfare gain. These exercises underscore the need for rich heterogeneity on the household side, as well as realistic restrictions on minimum house size and carefully calibrated supply and demand elasticities, if one is to arrive at a realistic quantification of the welfare effects.

The negative welfare cost of OOT demand prompts us to evaluate a tax on OOT investors. This exercise is calibrated to the experience of Vancouver which introduced such a tax in August 2016. We make two changes to the model to make this exercise more realistic. First, we make OOT buyers price-elastic and calibrate the elasticity to the observed response of OOT demand in Vancouver to the August 2016 tax. Second, the extra tax revenue funds additional public goods over which local residents derive utility. For the most plausible values of the public goods utility parameter, we find that a tax on OOT investors of the Vancouver magnitude offsets two-thirds of the welfare cost. Using the additional property tax revenue to lower property taxes on locals results in similar welfare gains from the OOT tax. The welfare gains from taxes flow mostly to renters when OOT demand is fairly price elastic, whereas they flow mostly to owners when demand is fairly price inelastic. Responding to the OOT inflow by instead relaxing the constraints on housing supply can also offset the welfare costs. This policy strongly favors renters.

Our paper contributes to three literatures. A large literature in finance solves partial-equilibrium models of portfolio choice between housing (extensive and intensive margin), financial assets, and mortgages. Examples are Campbell and Cocco (2003), Cocco (2005) and Yao and Zhang (2004) and Berger, Guerrieri, Lorenzoni, and Vavra (2017). Davis and Van Nieuwerburgh (2015) provides a recent summary of this literature. More recent work in macro-finance has solved such models in general equilibrium, adding aggregate risk, endogenizing house prices and sometimes also interest rates. E.g., Favilukis, Ludvigson, and Van Nieuwerburgh (2017) and Kaplan, Mitman, and Violante (2017). In this vein, Imrohoroglu, Matoba, and Tuzel (2016) study the effect of the passage of Proposition 13 which lowered property taxes in California and find quantitatively meaningful effects on house prices, moving rates, and welfare. 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.

Our model also connects to a growing empirical literature that studies the effect of OOT buyers on local housing markets. Badarinza and Ramadorai (2018) attribute foreign inflows in the London real estate market to political risk in the countries from which the capital flows originate. Using political shocks in a source country as an exogenous instrument, they estimate the effects of OOT buyers on house prices in London neighborhoods with a large pre-existing share of residents born in that source country. They find substantial price effects in such areas, which they interpret as safe haven effects. Sá (2017) also finds positive price effects on UK house prices and home ownership rates, using data on properties owned by overseas companies. Cvijanovic and Spaenjers (2018) use deeds records for Paris and find that non-resident foreigners typically buy in more desirable neighborhoods, pay more for these properties than locals and have lower capital gains upon a sale, and that their purchases are in part driven by local economic conditions in their country of origin. OOT buyers account for about 16.6% of purchases in Paris. They also show that 60% of OOT properties are used as secondary residence rather than being rented out as an investment property. More broadly, Bayer, Geissler, Magnum, and Roberts (2011), Chinco and Mayer (2016), and DeFusco, Nathanson, and Zwick (2017) all highlight the importance of investors in the dynamics of housing investment and price formation. Badarinza, Ramadorai, and Shimizu (2018) and Agarwal, Sing, and Wang (2018) emphasize the importance of the nationality of the parties in a commercial real estate transaction. We complement this literature by providing a model to confront the empirical evidence on OOT investors with. Such a model can be used for counter-factuals and policy analysis like in our Vancouver OOT tax exercise. We also provide new testable implications regarding the effects of OOT purchases on wages and inequality.

Finally, our work contributes to the urban economics literature which studies housing affordability. These models tend to be static and households tend to be risk neutral and/or have quasi-linear or CARA preferences.[1] The lack of risk, investment demand

[1] Hizmo (2015) and Ortalo-Magné and Prat (2016) study a portfolio choice problem where households

for housing, and wealth effects makes it hard to study the effect of OOT home investors on housing affordability in such a setting. Rather than assuming absentee landlords, we close the housing market by explicitly modeling local landlords. On the other hand, the urban models have an interesting spatial aspect, either within the city or across cities. 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). An earlier version of this paper considered a model with two neighborhoods in the city. The setup could accommodate spatial concentration of OOT demand in one neighborhood and had additional predictions for house price spillovers from an OOT shock to the other neighborhood and for between-neighborhood inequality. The aggregate welfare implications of OOT demand were very similar leading us to simplify the model along this dimension.

Because it is a heterogeneous-agent, incomplete-markets model, agents choices and equilibrium prices depend on the entire wealth distribution. We use state-of-the-art methods to solve the model, extending the approach of Favilukis et al. (2017), which itself extends Gomes and Michaelides (2008) and Krusell and Smith (1998) before that.

The rest of the paper is organized as follows. Section 2 provides stylized facts on the importance of OOT home buyers. Section 3 sets up the model. Section 4 calibrates the baseline model. Section 5 discusses its main results and discusses how the welfare results depend on parameter assumptions. Section 6 studies the effects of a tax on OOT investors, calibrated to the case of Vancouver. Section 7 concludes. The appendix contains further details on the model and on the data.

2 Motivating Evidence

Comprehensive data for OOT home buyer demand across time and space is not available, despite recent improvements in data collection. However, there is mounting evidence for the quantitative importance of OOT investors on both quantities and prices in the housing market of many major cities around the world. This section collects some stylized facts as motivation for (i) the theoretical model that is the focus of this paper, (ii) the calibration to a 10% OOT investor share, and (iii) the OOT tax policy exercise in Section 6.

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. 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).

New York We obtain data from *CoreLogic* on the fraction of OOT housing purchases for Manhattan and for the rest of the New York MSA between January 2004 and September 2016. Housing purchases are defined as purchases of single-family, 2-4 family, condominiums, and co-ops. OOT purchases are identified using the reported mailing addresses on payment/tax forms. Specifically, if the address of a buyer is either abroad or not contained in the list of 1,304 ZIP codes inside NYC MSA, then the transaction is classified as an OOT purchase. Appendix D.2 describes in detail how we deal with purchases made by LLCs and other corporate entities. Figure 1 plots the OOT purchase share at quarterly frequency, averaging the months in the quarter. Four facts stand out. First, OOT purchases are a non-trivial part of the market throughout the sample period. Second, the OOT share is much larger for Manhattan than for the rest of the metro area. Over the period of measurement, the data suggest an OOT share of 10% for Manhattan and 5% for the entire NYC metro area. Third, there is a steady increase in the OOT share in both Manhattan and the rest of the metro area over the last thirteen years. Fourth, Figure 1 shows that the OOT demand is quite persistent.

Figure 1: OOT Purchase Share New York

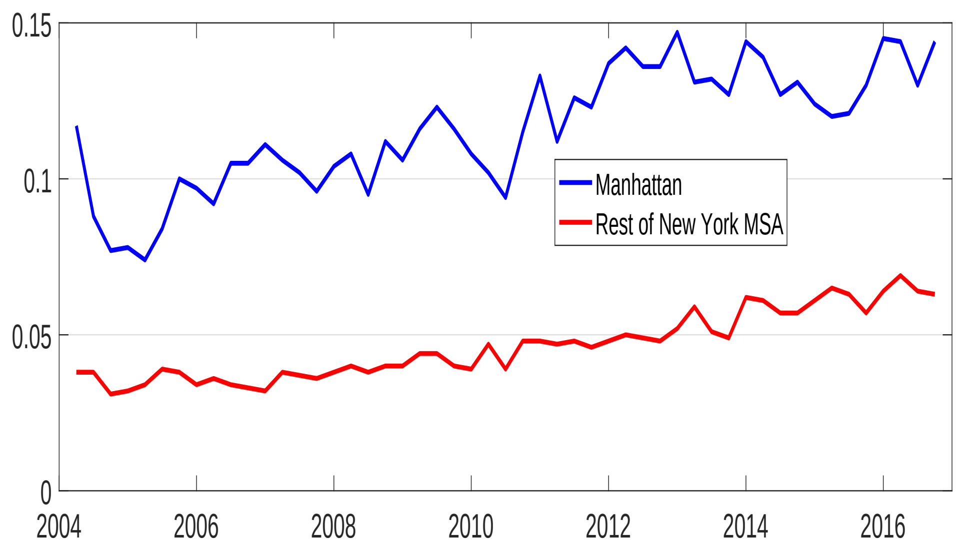

Suher (2016) provides corroborating evidence from New York tax records suggesting that of a sample of 84,000 condo units in the core of Manhattan in 2012, 25% are owned by non-residents.

Evidence for the prevalence of vacant properties owned by non-residents comes from the American Community Survey data. It reports homes that are vacant for seasonal, recreational or other occasional use. Data for Manhattan show that the number of such vacant units has risen from 30.9% of vacant units and 4.5% of occupied units in 2010 to 39.3% of vacant and 5.9% of occupied units in 2015. This category of vacancies under-counts non-resident ownership because the category can only be assigned if no one responds to multiple contact requests by the surveyor. At that point, the surveyor will speak to neighbors, brokers, or property managers to try to figure out whether the unit is in fact occupied but only used occasionally. Wheaton (2017) studies the ACS seasonal/occasional vacancy data and shows that they are related to regional variation in the amplitude of the house price boom and bust.

While non-resident buyers face no additional taxes to buy an apartment in NYC, despite public outcry for the introduction of such a tax, they face higher taxes than residents when selling. This includes Federal (20%) and New York State (8.8%) capital gains taxes as well as Federal estate taxes (no exemptions for foreigners). The Regional Plan Association backed a pied-à-terre tax for New York City in its 2018 Fourth Regional Plan, noting 60,000 NYC apartments are vacant but not on the housing market.

Rest of U.S. The National Association of Realtors (NAR) has been conducting a survey among its members on foreign home purchases in the United States since 2010. The share of foreign purchases was 10% during the 12-month period that ended March 2017 ($153bn). It was about 8% in the March 2015 ($103.9bn), March 2016 ($102.6bn), and the March 2018 ($121bn) surveys. The 2018 foreign investment is approximately double the 2010 investment ($65.9bn) when data collection started. Florida attracted most of the foreign investment in 2018 (22%). Foreigners accounted for 50% of purchases in Miami-Dade county. About 40% of foreign purchases are by non-resident foreigners. OOT purchases would exclude purchases by foreigners that are local residents but include all purchases by non-resident (out-of-MSA) U.S. citizens. Such data are not readily available for the entire U.S.

Chinco and Mayer (2016) use housing transactions merged with tax assessor data to identify OOT buyers, using the property tax billing address. They find that the OOT share rises as high as 17% in some boom markets like Las Vegas. Bayer et al. (2011) use transaction data for the period 1988-2009 for Los Angeles county. They focus on the role of all investors, without distinguishing between local and out-of-town investors. Using three different measures, the investor share triples between the early 1990 to the peak of the boom in 2003-06.[2]

[2]Their measures of investors are (i) whether an individual owns two homes at the same time –this measure rises to a nearly 30% share in 2006,– (ii) purchases that were resold within two years –15% of all homes bought in 2003-05 were resold within a two year period,– and (iii) flippers defined as the fraction of purchasers who buy at least two houses while holding them for less than two years –this measure peaks at 5%

London A study by the Mayoral office shows that 13% of properties sold in 2014-16 were bought by foreigners. Half of these homes were below £500,000, homes that could be bought by typical first-time home buyers. Separately, Knight Frank estimates that 49% of Central London homes above £1 million were sold to foreigners in the year to June 2013. More than half of these sales (28% of total sales) were to non-resident foreigners.

In April 2016, the U.K. began to levy a 3% stamp duty surcharge on second property purchases. Given the progressive schedule of stamp duty and the fact that OOT purchases are usually higher-end properties and always second homes, the incidence of this surcharge disproportionately falls on OOT buyers.

Paris Data from Cvijanovic and Spaenjers (2018) allow us to compute the share of OOT buyers in Paris. We define OOT buyers as either foreign or French residents who do not live in Paris or in Ile-de-France at the time of their housing purchase in Paris. These buyers account for 16.6% of transactions (and a somewhat larger share of value) over the 1992-2016 sample. The OOT purchase share in Paris is fairly constant over time.

The Paris data also contain information on the use of the apartment, distinguishing between primary residence, investment property for rent, and secondary home (pied-à-terre). This information is available for all sales, and sellers can be grouped by residence at the time of sale. This data shows that 61% of Parisian dwellings sold by OOT owners are secondary residences, with the remaining 39% for rent.

Vancouver and Toronto Foreigners owned 7.9% of Vancouver's and 7.2% of Toronto's condominium apartments in mid-2017 according to Statistics Canada. Non-residents owned 5.1% of all residential properties by value in Vancouver and 3.0% in Toronto. Statistics Canada suspects that these numbers understate foreign ownership because they fail to account for purchases made by corporations incorporated in Canada but controlled by foreign residents, as well as for properties bought by foreigners but in the name of a Canadian resident relative.[3]

Soaring prices prompted the government of British Columbia to introduce a 15% transaction tax on foreign buyers in August 2016. Foreigner buyers exclude foreign nationals who work and pay income taxes in British Columbia. This is consistent with the treatment of OOT buyers in our model: OOT investors do not work in the city either. The Government of Ontario followed suit with its own 15% transaction tax in April 2017. The

in 2006.

[3]See https://vancouversun.com/news/local-news/foreign-ownership-data-released-so-far-just-the-tip-of-the-iceberg-statscan-director.

levy in British Columbia was increased to 20% in mid-2018, and now applies to a larger geographical area. In the five weeks leading up to the introduction of the tax, foreign buyers accounted for 10% of purchase value in the Vancouver metro area and 11% in the City of Vancouver.

In April 2018, the City of Vancouver instated an "empty home tax" of 1% of the taxable assessed value of the home for properties that are left vacant, i.e., that do not serve as primary residence or are not rented out at least six months per year. The province of British Columbia passed its own "speculator" tax, which goes into effect in 2019, targeting both foreign and domestic OOT home buyers who don't live in or rent out their properties for at least six months per year. The BC vacancy tax rate is 2% of the home value for foreigners, 1% for OOT residents, and 0.5% for residents.

Australia and New Zealand Foreigners must apply with the government to purchase residential real estate in Australia. Data from Australia's Foreign Investment Review Board show a quadrupling of residential real estate approvals from 10,000 per year in the 2010-2012 period to 40,000 a year in 2015-2016. The 40,000 approvals correspond to $72 billion in investment. The foreign buyer share hovered between 10 and 15% between 2014.Q4 and 2017.Q1 according to the NAB.

Approvals dropped to 13,000 in 2016-2017, totalling $25 billion after the passage of stricter controls and higher application fees on foreign ownership. The province of New South Wales, where Sydney is located, introduced a stamp duty surcharge of 4% on the purchase of residential real estate by foreign persons from June 2016 onwards. For purchases after July 2017, the surcharge was increased to 8%. Victoria with capital Melbourne introduced its own stamp duty surcharge of 3% on July 1, 2015, and increased it to 7% for purchases after July 1, 2016. Foreigners who live and work in Australia are exempt. The introduction of a "ghost" tax on properties that are not available for rent or occupied more than half of the year and tighter Chinese capital controls may also have contributed to the decline in the foreign buyer share. By 2017.Q4, it had fallen to 8.4%.

New Zealand similarly passed severe limitations on foreign home buyers in August 2018, after house prices surged. The non-resident share of home purchases peaked in April 2018 at 3.3% in New Zealand, but at nearly 10% in Queenstown and 7.3% in Auckland.

Singapore and Hong Kong The non-resident foreign purchase share in Singapore was 7.4% in 2016.H1 and 6.8% in 2017.H2. Singapore imposed a special tax of 15% on foreign property buyers in 2013, and increased its tax to 20% in August 2018.

Hong Kong introduced a 15% non-resident stamp duty in 2012. Non-residents currently pay 30% stamp duty in Hong Kong.

Israel The share of residential properties owned by non-residents rose sharply from about 1% prior to 2000 to 6% in 2005-06 before falling back to 3% by 2009, according to the Israeli Ministry of Finance. The share is much higher in Tel Aviv (about 8-10%) and in Jerusalem (about 12-16%). Starting January 2017, the purchase tax for non-residents was substantially increased, relative to that for residents.

3 Model

We model a housing market, which we think of as corresponding to a metropolitan statistical area. The MSA has a fixed population normalized to one.[4]

3.1 Households

Preferences The economy consists of overlapping generations of households. There is a continuum of households of a given age. Each household maximizes utility u over consumption goods c, housing h, and labor supply n. Housing is divisible and there are no moving costs. Preferences are CRRA with risk aversion parameter γ and an aggregator function $\mathcal{C}$ defined below:

$$\begin{aligned} U(c_t, h_t, n_t) &= \frac{\mathcal{C}(c_t, h_t, l_t)^{1-\gamma}}{1-\gamma}, \\ n_t &= 1 - l_t. \end{aligned} \tag{1}$$

Total (non-sleeping) hours in a period of time are normalized to 1. This time is split between work n_t and leisure l_t. Since the earnings data and the model exclude the unemployed, we impose a minimum constraint on the number of hours worked.

There are two types of households in terms of the time discount factor. One group of households have a high degree of patience β^H while the rest have a low degree of

[4]Future work could consider an open-city model extension and study interactions between OOT investor flows and resident net migration patterns. Such a model would need to take a stance on a reservation utility of moving to other locations and on the moving costs. These would naturally differ by age, productivity, and wealth. A proliferation of parameters would ensue with little guidance from the literature on how to set these parameters. Evidence from the New York metropolitan area, discussed in Appendix D.3, suggests that net migration rates during the period of OOT inflows were small, and positively rather than negatively correlated with OOT purchase activity across time and space.

patience β^L. This preference heterogeneity helps the model match observed patterns of home ownership and wealth accumulation. 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, persistent labor productivity shock z.

There is an exogenous retirement age. After retirement, households earn a pension which is the product of an aggregate component $\overline{\Psi}$ and an idiosyncratic component ψ^z which has a cross-sectional mean of one. The idiosyncratic component reflects productivity during the last period of the working stage. Labor income is taxed at rate τ^{SS} to finance the pension system.

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. One interpretation is that attitudes towards saving are passed on from parents to children. 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 stylized facts: many households receive no bequest and there is substantial heterogeneity among bequest sizes for those who receive a bequest.[5]

Tenure Choice Let S_t be the aggregate state of the world, which includes the wage W_t, as well as the housing price P_t, rental cost R_t and previous housing stock H_{t-1}. 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 that there is one set of Bellman equations for each β-type. The household chooses each period whether to be an owner or a renter. Denote by V the value functions over these choices, with subscript R denoting a choice of

[5] 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, and that bequests arrive at different points in the life cycle for different households.

renting and O one of owning. The household solves:

$$V = \max \{V_O, V_R\}.$$

The Bellman equations for V_R and V_O are defined below.

Renter Problem A renter household chooses non-durable consumption c_t, housing consumption h_t, and working hours n_t to solve:

$$\begin{aligned}
&V_R(x_t, z_t, a, S_t) = \max_{c_t, h_t, n_t} U(c_t, h_t, n_t) + (1 - p^a)\beta E_t[V(x_{t+1}, z_{t+1}, a+1, S_{t+1})] \text{ s.t.} \\
&c_t + R_t h_t + Q b_{t+1} = y_t^a + \overline{\Psi}\psi^z + x_t \\
&x_{t+1} = b_{t+1} + \widehat{b}_{t+1} \geq 0 \\
&y_t^a = (1 - \tau^{SS}) W_t n_t G^a z_t \\
&n_t \geq \underline{n} \; if \; a < 65, \quad n_t = 0 \; if \; a \geq 65 \\
&h_t \geq \underline{h}
\end{aligned}$$

The renter's savings in the risk-free bond, Qb_{t+1}, where Q is the bond price, are obtained from the budget constraint. Next period's financial wealth consists of these savings plus any accidental bequest received, $\widehat{b}$. Retirement income $\overline{\Psi}\psi^z$ is zero prior to age 65 and labor income is zero after age 65. There is a minimum dwelling size $\underline{h}$ and a minimum number of hours worked $\underline{n}$.

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

$$\begin{aligned}
&V_O(x_t, z_t, a, S_t) = \max_{c_t, h_t, \widehat{h}_t, n_t} U(c_t, h_t, n_t) + (1 - p^a)\beta E_t[V(x_{t+1}, z_{t+1}, a+1, S_{t+1})] \text{ s.t.} \\
&c_t + P_t h_t + Q b_{t+1} + P_t \widehat{h}_t = y_t^a + \overline{\Psi}\psi^z + x_t + R_t \widehat{h}_t \\
&x_{t+1} = b_{t+1} + \widehat{b}_{t+1} + P_{t+1}(h_t + \widehat{h}_t)(1 - \delta - \tau^P) \\
&-Q_t b_{t+1} \leq P_t(\theta_{res} h_t + \theta_{inv} \widehat{h}_t) - R_t \widehat{h}_t - (y_t - c_t) \\
&n_t \geq \underline{n} \; if \; a < 65, \quad n_t = 0 \; if \; a \geq 65 \\
&h_t \geq \underline{h} \\
&\widehat{h}_t \geq 0
\end{aligned}$$

Local home owners are the landlords to the local renters. For simplicity, we assume that renters cannot buy investment property. Owners earn rental income $R_t \widehat{h}_t$ on their investment units.

The physical rate of depreciation for all housing units is δ. The term $Ph\delta$ is a financial costs, i.e., a maintenance cost.[6] As shown in equation (6) below, the physical depreciation $Ph\delta$ can be replaced by residential investment undertaken by the construction sector. This treatment of depreciation avoids having to keep track of the aggregate owner-occupied fraction of housing as an additional state variable.

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. In the baseline model, local government spending confers no utility. In the model of Section 6, property tax revenue finances a local public good that enters the local residents' utility function.

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 $\theta_{res} \geq \theta_{inv}$. Rental income and savings $(y - c)$ in the first period (four years in the calibration) after the purchase of a rental property do not count towards a household's net worth when computing the LTV ratio.

3.2 Firms

Goods Producers There are a large number n_f of identical, competitive firms all of which produce the numéraire consumption good.[7] 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:

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

These firms are owned by national equity owners, and all profits are distributed outside of the city. Below, we explore sensitivity to this assumption.

Developers There are a large number n_f of identical, competitive construction firms which produce new housing units and sell them locally at a price P_t per unit. Like the

[6]It is easy to solve for a model where investment housing incurs an additional maintenance cost which is a fraction $\delta_{inv} \geq 0$ of the value of the property. Increasing δ_{inv} makes renting less attractive and increases the home ownership rate. We do not need a depreciation wedge to generate the observed home ownership rate.

[7]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.

consumption good firms, construction firms are owned by national equity owners, and all profits are distributed outside of the city. Let H_{t-1} be the existing housing stock. The construction firms have decreasing returns to scale and choose labor to maximize profit each period:

$$\Pi_t = \max_{N_{h,t}} \left(1 - \frac{H_{t-1}}{\overline{H}}\right) P_t N_{h,t}^{\rho_h} - W_t N_{h,t} \tag{3}$$

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 parameter $\overline{H}$ is an upper bound on the total housing that can be built in the city. We interpret $\overline{H}$ as the total land area zoned for residential use multiplied by the maximum permitted number of floors that can be built on this land, the floor area ratio (FAR). The term $1 - \frac{H_{t-1}}{\overline{H}}$ captures the idea that, the more housing is already built in a location, the more expensive it is to build more. For example, additional housing may have to take the form of taller structures, buildings on less suitable terrain, or irregular infill lots. When $\overline{H}$ is sufficiently high, the model's solution becomes independent of $\overline{H}$, and the housing supply elasticity is governed solely by ρ_h. When $\overline{H}$ is sufficiently low, the housing supply elasticity is lower and depends on both $\overline{H}$ and ρ_h.[8]

3.3 Out-of-town Buyers

We assume that out-of-town (OOT) buyers inelastically demand some amount of housing. In Section 6, we relax this assumption and introduce a non-zero elasticity of OOT demand to house prices. OOT home buyer demand is stochastic and is the only source of aggregate risk in the model. One can think of shocks to the oil price, shocks to the exchange rate (Ruf and Levi, 2011), or political shocks in the country of origin (Badarinza and Ramadorai, 2018) or destination country (e.g., Brexit) as some of the underlying drivers of this stochastic OOT process.

OOT demand follows a 2-state Markov process with a low state $OOT_t = OOT^L$ and a high state $OOT_t = OOT^H > OOT^L$. We assume a symmetric transition probability matrix with the parameter π governing the expected duration of each regime.

A key assumption is that housing owned by OOT buyers is not rented out to locals. Section 2, and in particular Cvijanovic and Spaenjers (2018) provide empirical support for this assumption. Below, we solve a version of the model that assumes OOT housing is rented out, with substantially different results. OOT buyers do not work in the local labor

[8]In this sense, the model captures that construction firms must pay more for land when land is scarce. This scarcity is reflected in equilibrium house prices.

market, do not consume local public goods, and are not counted as part of city welfare.

3.4 Equilibrium

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

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

$$N_{c,t} = \left(\frac{\rho_c}{W_t}\right)^{\frac{1}{1-\rho_c}} \quad \text{and} \quad N_{h,t} = \left(\frac{\left(1 - \frac{H_{t-1}}{\overline{H}}\right) P_t \rho_h}{W_t}\right)^{\frac{1}{1-\rho_h}}. \tag{4}$$

Second, labor markets clear:

$$n_f \left(N_{c,t} + N_{h,t}\right) = N_t. \tag{5}$$

Third, the housing market clears:

$$(1-\delta)H_{t-1} + \underbrace{n_f \left(1 - \frac{H_{t-1}}{\overline{H}}\right) N_{h,t}^{\rho_h}}_{INV} = \widehat{H}_t + H_t^O + OOT_t. \tag{6}$$

The left-hand-side is the supply of housing which consists of the non-depreciated housing stock and residential investment INV. The right-hand-side is the demand for those housing units by owner-occupiers and landlords. Fourth, the rental market clears:

$$\widehat{H}_t = H_t^r \tag{7}$$

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

$$\overline{\Psi} N_{ret} = \tau^{SS} E\left[N_t W_t\right], \tag{8}$$

where we used the fact that G_a and z average to 1 in the cross-section, and N_{ret} is the total number of retirees, which is a constant.[9] Sixth, the aggregate state S_t evolves according

[9]For simplicity, we assume that the total pension payments are equal to the average of all social security tax revenues, where the expectation is across high and low OOT demand states. OOT demand affect wages and therefore the total social security tax collected in a city. We do not think that letting the pension fluctuate

to rational expectations. Seventh, the value of all bequests received is equal to the wealth of all agents who die. Appendix A presents first order conditions.

3.5 Welfare effects on Locals from OOT Buyers

We compute the welfare effect of OOT home buyers using the following procedure. Suppose that OOT demand in period t is low and that it stays low in period $t+1$. Denote agent i's welfare at $t+1$ as $V_{t+1,i}(LL)$. Suppose instead that OOT demand at $t+1$ switches to high and denote agent i's welfare in this situation as $V_{t+1,i}(LH)$. Agent i would be willing to give up $\mathcal{W}_{t,i}$ in consumption equivalent units to stay in the low foreign demand state, where:

$$\mathcal{W}_{t,i} = -1 + \left(\frac{V_{t+1,i}(LL)}{V_{t+1,i}(LH)} \right)^{\frac{1}{(1-\gamma)(1-\alpha_n)}} .$$

We compute aggregate welfare effects from "inflows" by summing $\mathcal{W}_{t,i}$ across all agents, calling the resulting aggregate welfare measure $\mathcal{W}$. We also sum separately among owners ($\mathcal{W}_o$) and renters ($\mathcal{W}_r$), for different age groups, income groups, and wealth groups.

4 Calibration

The baseline model is calibrated to the average U.S. metropolitan area. Table 1 summarizes the chosen model parameters. Externally calibrated parameters are indicated with a star. After we discuss the baseline results, we switch off several features of the model, one by one, to quantify the role of the various model ingredients. Section 6 discusses a calibration to Vancouver with elastic OOT demand and utility from public goods.

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. People age 65 and over comprise 19% of the MSA population age 21 and over in the data (see appendix D.1) and 21.8% in the model.[10]

with OOT demand of local real estate would be desirable. In the U.S., Social Security is maintained at the national level, and pension payments do not depend on local-area variation in wages.

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

Table 1: Calibration

Description	Parameter	Values	Target
Panel A: **Labor Income and Pension**			
Income states	z	[0.258 0.775 2.193]	Income distribution SCF
Pension income*	ψ^z	[0.484 0.968 1.582]	Pension distr. rules
Pension tax rate*	τ^{SS}	0.10	Contribution rates
Panel B: **Preferences**			
Time Preference (4yr)	(β^L, β^H)	(0.799,1.003)	Wealth/income and wealth gini
Risk aversion*	γ	5	Standard value
Leisure weight*	α_n	0.66	Average workweek
Housing consumption weight*	α_h	0.32	Housing share in consumption of 21.6%
CES parameter consumption/housing*	ϵ	-0.5	Intra-temp elasticity c & h of 2/3
CES parameter consumption/leisure	η	0	Frish elasticity of labor supply of 1
Minimum hours	$\underline{n}$	0.5 E[n]	Capture part-time and full-time work
Panel C: **Finance and Regulatory**			
Bond Price (4yr)*	Q	0.8213	Avg return stock-bond portfolio 1987-2018
Maximum residential LTV*	θ_{res}	0.90	Max LTV on primary residence
Maximum investment LTV*	θ_{inv}	0.80	Max LTV on investment property
Property tax (4yr)*	τ^P	4%	Property tax rates
Housing depreciation (4yr)*	δ	0.094568	BEA housing depr 1987-2018
Minimum dwelling size*	$\underline{h}$	$1/3\ E[h]$	Annual Housing Survey
Panel D: **Production and OOT Demand**			
Return to scale*	$\rho_c = \rho_h$	0.66	Labor share of income
Available space	$\overline{H}$	3.0	Housing supply elasticity of 1.0
OOT demand*	(OOT^L, OOT^H)	(0.00,0.10)	Section 2
OOT demand transition prob.*	π	0.90	Section 2

Notes: The table reports the parameters of the model, their values in the baseline model and the the Vancouver model, as well as the target. Moments denoted with a * are externally calibrated. The other parameters are chosen to target a moment that is endogenous to the model.

Labor Income Pre-tax labor earnings 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. The choice of hours is subject to a minimum hours constraint, which is set to 0.5 times average hours worked. This constraint rules out a choice of a positive but very small number of hours, which we do not see in the data given the indivisibility of jobs. It also rules out unemployment since the earnings data is for the (part-time and full-time) employed. This constraint binds for 9% of workers in equilibrium. The presence of the constraint enables the model to better match the observed correlation between wealth and income.

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 the mean of labor earnings by age. We use data from ten waves of the Survey of Consumer Finance (SCF) from 1983-2010 to estimate G^a.

The idiosyncratic productivity process z^i is chosen to match earnings inequality and to generate realistic persistence in individual earnings. We discretize productivity z as a 3-state Markov chain. The values for the three states are chosen so that the average earnings of households in the bottom 25%, middle 50%, and top 25% of the earnings distribution in the model matches those same objects in the SCF. We assume a parsimonious transition probability matrix for z, where the probability of staying in the same productivity state is 90% for the average worker and 100% for retirees. The implied standard deviation and autocorrelation of the idiosyncratic component of labor income are 0.7 and 0.9, respectively, matching the micro evidence in Storesletten, Telmer, and Yaron (2006). We refer the reader to appendix B for details.

The Social Security tax rate,τ^{SS}, is set to 10%. In the data, employees contribute 6% and employers contribute an additional 6%, but only on income below \$118,500. 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, resulting in the vector ψ^z. The average pension $\overline{\Psi}$ is determined by equation (8) to balance the social security budget.

Preferences The functional form for the utility function is given in equation (1). We set risk aversion $\gamma = 5$, a standard value in the macro-finance literature. We assume the CES aggregator:

$$\mathcal{C}(c_t, h_t, l_t) = \left[(1-\alpha_n)\left((1-\alpha_h)c_t^{\epsilon} + \alpha_h h_t^{\epsilon}\right)^{\frac{\eta}{\epsilon}} + \alpha_n l_t^{\eta}\right]^{\frac{1}{\eta}}$$

We set α_n to match the average workweek of 34.5 hours per week in the U.S., which is 30.8% of non-sleeping time (112 hours per week). The model generates 31.8%. We set α_h to match the ratio of housing consumption to total consumption. The average in U.S. data for 1987-2018 (BEA real housing services consumption to real non-durable plus real services consumption) is 22.1%; the model generates 22.9%. The parameter ϵ governs the intra-temporal marginal rate of substitution between housing and non-housing consumption. While the literature contains a range of estimates, the median elasticity across studies is 2/3, resulting in $\epsilon = -0.5$. The parameter η governs the Frish elasticity of labor supply. (For our utility function, the Frisch elasticity depends not only on η but also on the other utility parameters.) Here too, there is a wide range of estimates from 0 to 1.5 based on micro data and from 2 to 4 based on macro data. We target a value of 1, resulting in $\eta = 0$. We obtain a Frisch elasticity in the model of 1.01 using aggregate quantities and 0.98 when using individual quantities.

We set $\beta^H = 1.003$ and $\beta^L = 0.799$. Note that because of the probability of death, the effective discount rates are $(1 - p(a))\beta$. A 25% share of agents has β^H, the rest has β^L. This delivers an average β of 0.85 (0.96 per year), set to match the average wealth-income ratio, which is 5.66 in the model and 5.69 in the 1998-2010 SCF data. The dispersion in betas helps the model generate a wealth Gini coefficient of 0.78, matching the 0.78 in 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, 1.4% of wealth is bequeathed each year, close to the 1.2% in the data.

Geography and Production The return to scale parameter in the both sectors is set to $\rho_c = \rho_h = 0.66$ to match the observed labor share of output of 66%.

Given the returns to scale parameter in the construction sector ρ_h, the available residential land $\overline{H}$ determines the long-run housing supply elasticity. Appendix C contains the derivation. Saiz (2010) estimates housing supply elasticities for the largest 50 MSAs; they range from 0.6 for Miami to 1.70 for Albany. We target a value of 1.0, which is the elasticity of Tampa, the 20^{th} largest one among the 50 MSAs. It is also the weighted average of the housing supply elasticities, weighted by the MSA population shares. The model generates a supply elasticity of 0.97.

Housing 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 $\tau^P = 0.04$ or 1% per year. This is the median property tax rate among U.S. states in 2017.

We assume that property depreciates at 2.45% per year and set $\delta = 0.09457$. This is the average depreciation rate for privately-held residential property in the BEA Fixed Asset tables for the period 1987-2016.

Households in the model only have access to a real bond to save in (beside housing). In reality, a wider range of investment opportunities is available and governs wealth accumulation in the data. We set $Q = 0.8213$ to target the 5.05% average real annual return on a 50-50 bond-stock portfolio over the period 1987-2018. The bond is the five-year constant maturity Treasury and the stock is the CRSP value-weighted stock index.

We set the minimum housing size $\underline{h}$ to $1/3$ of the average housing size. According to the American Housing Survey, the average housing unit (including single- and multi-family units) built after 1960 is about 1,600 square feet big. The minimum house size then corresponds to 533 square feet.[11]

OOT share We set the low OOT state to reflect a situation without (before) OOT real estate purchases: $OOT^L = 0$. Given the empirical evidence discussed in section 2 for several of the world's major MSAs, we set the high state to $OOT^H = 0.10$, so that OOT ownership represents 10% of the housing demand. We assume a state transition probability of 10% per 4-year period. That is, $\pi = 0.90$. Given the uncertainty regarding these numbers, we explore changing both OOT^H and π below.

5 Baseline Model Results

5.1 Earnings, Wealth, and Home Ownership

A first check on the model concerns its ability to broadly match observed patterns in income, wealth accumulation and home ownership over the life cycle.

The top panels of Figure 2 show pre-tax earnings by age in the data (left panel) and in the model (right panel). The solid black line shows the well-known hump-shaped labor

[11] For example, New York City prohibited housing units smaller than 400 square feet until recently. We normalize the average house size in the model to be 1600 square feet.

income profile over the life cycle for the median household. The dashed red line shows average labor income in the bottom 25 percent of the labor income distribution, the dash-dotted blue line reports average income among the middle 50 percent of the income distribution, and the dotted green line shows the average income among the top-25% of the distribution. The model matches the observed labor income profiles closely. Since hours worked are a choice variable, the good fit of labor earnings means that model has realistic implications for hours worked. Our assumption of a constant income in retirement after age 65 causes a more discrete decline in income than in the data, presumably because some people continue to work after age 65 in the data.

While income is fairly directly pinned down by the productivity calibration, wealth and home ownership result from households' optimal consumption and investment decisions. The middle panels of Figure 2 plots household net worth; the data is again on the left (in 2010 real dollars) and the model is on the right. Net worth is defined as real estate wealth plus financial wealth minus debt. The bottom panels plot the home ownership rate.

The model generates a home ownership rate of 62.2% (and 61.05% prior to the OOT inflow), close to the observed home ownership rate of 60.6% for the average U.S. MSA in 2016. The model fits the life-cycle patterns of wealth accumulation and home ownership well. The average home ownership rate starts out below 20% for the youngest households and displays a hump-shaped pattern over the life-cycle. It peaks at about 80% around age 70 in both model and data. It then declines in retirement.

The model also generates about the right amount of average wealth at different ages during the working stage of life. Households accumulate about $250,000 by age 40 and $650,000 by age 65, on average, in both model and data. Wealth gradually declines in retirement, in large part because home ownership rate declines. Both the decline in home ownership and total wealth are steeper in the model than in the data, and closely connected to each other.[12]

The model generates substantial cross-sectional variation in wealth and home ownership across cohorts and income groups that is broadly consistent with the data. The model matches the Gini coefficient for wealth (.78 in model and .78 in data), which substantially exceeds the Gini of income (0.46 in model and 0.51 in data). The average wealth held by the top-25% of income earners is somewhat higher than in the data while the share of

[12] Allowing for an intentional bequest motive and/or adding late-in-life medical/long-term care risk would give households additional motives to slow down wealth decumulation (Ameriks, Caplin, Laufer, and Van Nieuwerburgh, 2011). So would giving elderly people a preference for "aging in place" or letting them forgo housing maintenance (Cocco and Lopes, 2017). Adding these motives would overly complicate the model whose main purpose is to analyze the effect of OOT investors.

Figure 2: Earnings, Net Worth, and Home Ownership across Age and Income groups

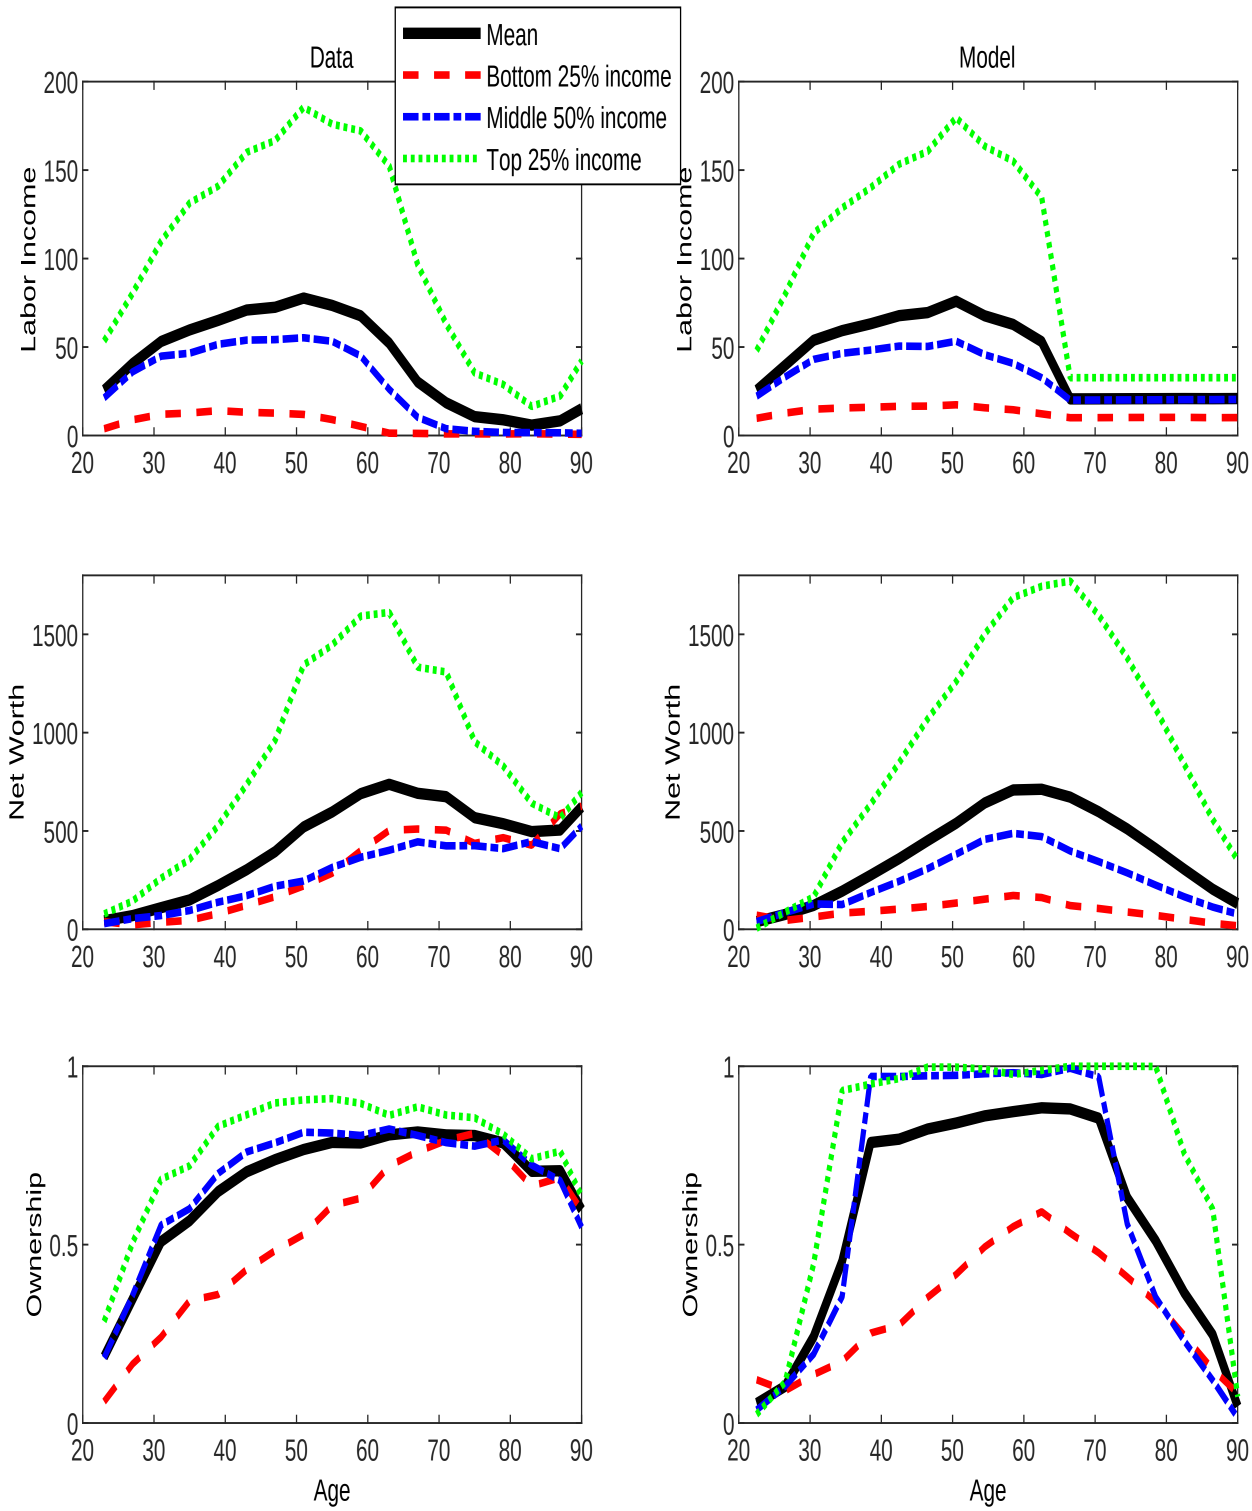

Notes: The left panels are for the data and based on the Survey of Consumer Finance (all 1983-2010 waves). The right panels are for the benchmark model. The top row denotes household earnings. The middle row denotes household wealth. The bottom row denotes the home ownership rate.

wealth held by this group is somewhat lower (61% in the data and 48% in the model).

5.2 House Prices and Rents

House prices, rents, and wages are determined in equilibrium. Since risk premia are small, as can be expected form a model with CRRA preferences and risk aversion of 5, the house price-to-rent ratio is close to the one arising from the user cost formula: $(1 - Q \times (1 - \delta - \tau^P))^{-1} = 3.46$ or 13.8 for the annualized price-rent ratio. This 13.8 ratio in the model is close to the 13.0 average across the 75 largest MSAs in the U.S. data (see Appendix D). The quantity and price results suggest that the model is well positioned to quantitatively evaluate both the average and the distributional consequences of OOT purchases.

5.3 Main Results: Increase in OOT Purchases

We now study the effect of an increase in OOT home buyer demand. Recall that OOT demand takes on two values, 10% of housing in the high state and zero in the low state, and is the only source of aggregate risk in the economy. Conditional on a switch, each regime is expected to last 40 years.

Table 2 summarizes how key prices and quantities adjust in response to the OOT shock. The first row reports the short-run response, in the first (four-year) period after the economy switches from the low to the high OOT state. These short-run changes are denoted by $\Delta_1 x$ for variable x, and expressed as percentage changes. The second row reports long-run changes, measured as the difference between the high-OOT and low-OOT stochastic steady states. These long-run changes are denoted by $\Delta_{ss} x$, and also expressed as percentage changes. The last three columns report welfare effects for the average renter, for the average home owner, and average across all households. They are also expressed as percentage changes and measure how much the household's value function (which incorporates the entire expected present discounted value of future changes) changes upon a transition from the low to the high OOT state.

The increase in OOT demand for housing is partly met by new construction. Residential investment increases by 17.33% in the short-run and by 9.58% in the long-run. Because of decreasing returns to scale ($\rho_h < 1$) and limited land $\overline{H}$, housing supply is not perfectly elastic even in the long run. Eventually, the housing stock increases by only 27.5% of the additional OOT demand.

To clear housing markets, locals must therefore consume less housing. These effects are stronger in the short-run since it takes time to build. Rents rise by 18.77% in the short-run to induce a sufficiently large decline in local housing demand. House prices rise by 10.05%. Thus, the arrival of OOT investors leads to a significant increase in the cost of housing in the first four years. It not only increases the prices of owner-occupied units, which the OOT investors buy, but it has an even larger effects on rents. In the long-run the cost of rental housing increases by 13.43% and the cost of owner-occupied housing by 7.06%.

The increase in rents exceeds that in price, so that the price-rent ratio falls modestly from 14.24 to 13.21 (-7.34%) on impact and rebounds to 13.43 in the long run (-5.68%). What explains the decline in the P/R ratio? Because the OOT demand process is mean-reverting, a high OOT state today implies lower future OOT demand and lower expected rental growth. The decline in the price-rent ratio reflects that lower expected future demand, i.e., a cash flow channel. The risk premium associated with the OOT demand is small, and does not fluctuate much with the OOT state. Discount rate variation explains little of the price-rent dynamics.[13]

The price-income ratio in the model, denoted HP/Y in the table, is computed as the average house price per square foot, P, multiplied by the average house size in square feet divided by the average earnings of working-age households Y. Because we want to compare the same house across OOT states, we hold the average house size in a zone fixed at its unconditional average. The price-income ratio increases substantially from 3.41 to 3.75, an increase of 9.97%. A second major housing affordability indicator is the fraction of income spent on rent, denoted HR/Y. Average rent is the rent per square foot R multiplied by the average house size divided by average earnings. The rent-income ratio rises by 18.81%. By these common metrics, OOT buyers reduce housing affordability substantially.

The home ownership rate, Own, increases by 5.01% (from 61.05% to 64.10%) when OOT buyers first enter and by 5.22% in the long-run. The home ownership rate includes the OOT buyers, consistent with the data. We assume each OOT buyer purchases a housing unit of the same size as the average local home owner. One might conjecture that OOT buyers displace local owners since the housing stock does not expand in proportion to OOT housing demand. However, the home ownership rate among local residents remains constant. The reasons are that (i) the decline in the P/R ratio makes ownership more attractive relative to renting for households that are affected by the downpayment constraint, and (ii) local owners adjust also on the intensive margin by living in smaller

[13]To generate meaningful variation in housing risk premia would require, additional sources of aggregate risk such as changes in mortgage lending standards as in Favilukis et al. (2017).

housing units (-9.07% in the short-run and -6.64% in the long run).

OOT demand drives up wages by 1.07% in he short-run and by 0.80% in the long-run. This happens because a higher housing stock requires more workers in the housing sector. After a positive OOT shock, the share of construction employment increases by 19.32% (construction share of hours worked rises from 6.47% to 7.72%) as additional housing is being built. It then slowly falls, but stays at a higher level as long as OOT demand remains high. This is because the higher housing stock requires more maintenance. The higher demand for labor in the housing sector begets higher wages because it siphons employees from the non-housing sector. The non-housing sector has a lower labor-to-capital ratio which pushes up wages. Higher city-wide wages lead to lower aggregate labor demand, with hours worked being 1.13% lower when OOT demand is high. In sum, OOT buyers prompt an increase in wages and a boom in the non-tradable sector, but also to a loss in competitiveness and a fall in employment of the tradable sector.[14] The increase in wages is offset by a reduction in hours worked so that average earnings are approximately unchanged by the OOT shock. Cost of housing increases translate one-for-one into cost of living increases.

Welfare What are the welfare effects from an increase in OOT demand? The average household in the baseline model is worse off from OOT buyers of local real estate, with a welfare loss of 0.74% in consumption equivalent units, as indicated in the column labeled $\mathcal{W}$.

To understand the headline welfare number, it is important to consider how different types of agents are affected. The average renter is severely hurt by the inflow and would be willing to give up 2.42% of lifetime consumption to avoid the OOT capital inflow. In contrast, an average home owner benefits by 0.56% of lifetime consumption.

The top panel of Figure 3 splits out these welfare gains and losses for owners and renters by age group. Renters are unequivocally hurt by the higher OOT demand as their current and future housing expenditures rise. Working-age renters receive some relief in the form of higher current and future wages. This benefits the young more, given their longer expected work life. An average 75-year old renter would be willing to give up 8.23% of lifetime consumption to avoid the OOT home demand, while an average 35-year old renter would give up only 1.39% of lifetime consumption. However, rents go up far more than wages.

[14] Although housing is the only non-tradable good in our model, the same intuition carries over to any other non-tradable goods and services that OOT consume when they are in town, e.g., restaurants or entertainment.

Table 2: Effect of OOT Demand

	W	R	P	INV	P/R	HP/Y	HR/Y	Own	$\mathcal{W}$	$\mathcal{W}_r$	$\mathcal{W}_o$
	Panel A: **Baseline Model**										
Δ_1	1.07	18.77	10.05	17.33	−7.34	10.05	18.77	5.01	−0.74	−2.43	0.56
Δ_{ss}	0.80	13.58	7.12	9.60	−5.68	7.34	13.80	5.22	–	–	–
	Panel B: **Size OOT**										
Δ_1 OOT^H=15%	1.69	29.41	15.50	26.83	−10.75	15.46	29.37	7.49	−1.19	−3.80	0.77
Δ_1 OOT^H=5%	0.48	8.97	4.90	8.42	−3.73	4.85	8.92	1.69	−0.32	−1.13	0.29
	Panel C: **Persistence OOT**										
Δ_1 π=.75	0.91	18.82	7.14	11.97	−9.83	7.47	19.18	4.92	−0.70	−2.03	0.27
Δ_1 π=.95	1.24	18.52	11.52	19.88	−5.90	11.49	18.49	5.52	−0.81	−2.69	0.65
	Panel D: **OOT rent out property**										
Δ_1	−0.06	0.07	0.02	0.15	−0.05	−0.11	−0.06	2.39	0.00	0.01	−0.00
	Panel E: **No patience heteroegeneity**										
Δ_1	1.06	18.51	9.94	17.08	−7.23	9.97	18.54	2.60	−0.37	−1.76	0.33
	Panel F: **LTV limit**										
Δ_1 LTV=.80	1.06	18.80	10.06	17.30	−7.35	10.03	18.76	6.64	−0.79	−2.38	0.63
Δ_1 LTV=.95	1.08	18.74	10.04	17.25	−7.33	10.05	18.75	5.09	−0.72	−2.52	0.50
	Panel G: **Housing supply elasticity**										
Δ_1 $\overline{H}$=1	1.18	19.59	10.76	17.17	−7.38	10.60	19.41	4.60	−1.24	−3.80	0.46
Δ_1 $\overline{H}$=1000	1.01	18.16	9.61	17.20	−7.23	9.68	18.23	5.54	−0.56	−1.95	0.54
	Panel H: **Housing demand elasticity**										
Δ_1 $(1-\epsilon)^{-1}$=0.33	1.66	30.61	15.24	26.30	−11.76	15.36	30.74	3.16	−1.08	−3.62	0.71
Δ_1 $(1-\epsilon)^{-1}$=0.5	1.32	24.02	12.36	21.43	−9.40	12.40	24.07	4.12	−0.92	−2.98	0.60
Δ_1 $(1-\epsilon)^{-1}$=1	0.75	13.44	7.56	13.04	−5.19	7.47	13.36	5.09	−0.56	−1.80	0.41
	Panel I: **No minimum house size**										
Δ_1	1.03	17.37	9.44	16.20	−6.75	9.54	17.47	6.43	−0.38	−1.43	0.51
	Panel J: **Profit redistribution**										
Δ_1	1.15	17.87	9.66	16.38	−6.96	9.50	17.69	9.02	−0.72	−1.80	0.75
	Panel K: **Only owners**										
Δ_1	1.14	17.38	9.45	16.03	−6.75	9.64	17.58	0.00	0.04	−−	0.04

Notes: The table reports averages from a long simulation of the model. It reports the wage W, the rent per sqft R, the house price per sqft P, the price-rent ratio (each expressed per square foot and rent expressed annually), the home value-to-income ratio HP/Y (price times average house size among all residents divided by average annualized income in the working-stage of life), the rental expenditure share (rent times average house size among all residents divided by annualized income in the working-stage of life), the home ownership rate *Own* (including OOT owners), the average welfare change from an increase in OOT buyers among all residents, $\mathcal{W}$, the welfare change among renters $\mathcal{W}_r$, and among owners $\mathcal{W}_o$. Panel A is for the baseline calibration. The first row, labeled Δ_1, shows the percentage change in the variables in the first period of the transition from the low to the high OOT housing demand state. The second row, labeled Δ_{ss} shows the percentage difference between the variables in the stochastic stead states of high and low OOT home demand. Panels B-K report variants of the baseline model, described in the main text.

Owners benefit from higher foreign demand because they reap capital gains on their house. This increases their wealth and expands their consumption opportunity set. Higher real estate prices induce local owners to consume less real estate and more other goods. This creates a welfare gain for owners who expect to remain owners but do not expect their housing consumption to rise. These households are not worried about the higher future rents brought about by the OOT inflow. The capital gains effect is weaker for younger owners because they tend to own relatively little real estate and expect their housing consumption to increase in the future as they move up the housing ladder. Some young owners even experience welfare losses. The positive wage effect we described for renters also applies to owners. The confluence of these factors generates a welfare benefit for the average owner that increases from 0.45% at age 25 to 0.69% at age 55. It then falls to +0.17% at age 70, because future rental risk starts to matter. At advanced ages, this concern recedes due to low remaining life expectancy and the capital gain effect dominates once more, resulting in larger welfare gains for the very old.

For owners, OOT inflows are not unambiguously good. About 12% of 25-year old owners are hurt by the inflow. This fraction goes down to about 10% for older owners. These are households whose bequest realizations were low and whose past income realizations were low until recently. They have low housing wealth relative to their current income. The expect their housing consumption to grow but have not yet accumulated much housing wealth. This limits their capital gain from inflows.

This discussion makes clear why the model needs to feature a rich cross-section of agents in terms of age, earnings, wealth, and home ownership. House price, rent, and wage changes induced by an increase in OOT demand will affect these agents differently. Panel B of Figure 3 reports the welfare effects for households sorted by income, panel C for households sorted by net worth. The welfare costs are disproportionately born by low-income households, especially younger and elderly low-income households. Not only are these households much more likely to be renters, they also have high marginal utility of consumption which makes them vulnerable to the increase in the cost of living caused by the OOT inflow. The welfare effects along the wealth distribution are similar, with elderly low-wealth households suffering the most and middle-aged, middle-class households benefiting the most.

Political Economy of OOT Purchases Although the average welfare effect of OOT buyers is negative, it is not evenly distributed. Owners make up 61% of the resident population prior to the inflow, and owners on average benefit from an inflow. Renters, which are a numerical minority in the model and in the average U.S. metro area, are hurt by

Figure 3: Welfare Effects of Increase in OOT Housing Demand

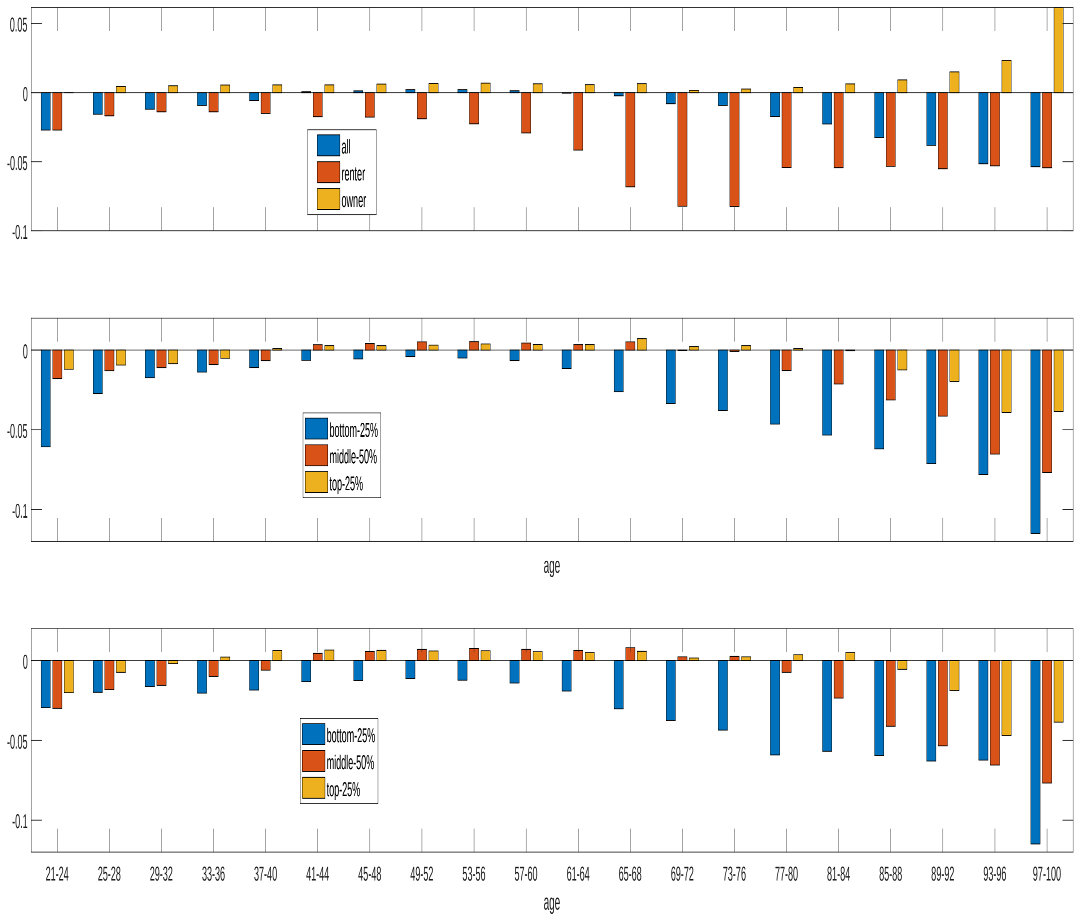

Notes: The figure presents the consumption-equivalent welfare change from an increase in OOT housing demand in the baseline model. All numbers are in percent, thus -0.05 for a group of households means that the average household in that group would need to receive 5% of its average consumption to be as well off after the OOT increase as before. All panels split households into 20 age groups of 4-years. The top panel splits households by tenure status, the middle panel by income (incl. retirement income), and the bottom panel by wealth (financial wealth, positive or negative, plus housing wealth).

much more than the owners gain. If each person received one vote, a majority of 55% of the population would prefer the OOT inflow to occur. However, these preferences are diametrically opposed for owners and renters: 94% of owners prefer the inflow, but only 4% of renters do. This calculation suggests that policies aimed at curbing OOT purchases not only have redistributive consequences, they may also be politically unpopular.

5.4 Exploring Key Drivers of the Model

To better understand the workings of the model and to gauge the sensitivity of its welfare implications, we now explore various versions of the model where we switch off model ingredients or change key parameters. Some key experiments involve the OOT housing demand process itself. Panels B-K of Table 2 report the results.

5.4.1 Size OOT Demand

The size of the OOT housing demand shock inflow matters. We study a case where the foreign inflow is 50% larger than in the baseline and one where it is 50% smaller. The welfare cost to society is 1.19% in the former case, and 0.32% in the latter case, straddling the 0.74% benchmark cost. The first 5% point increase in OOT demand from 5% to 10% increases the welfare cost by 84%. The second 5% point increase from 10% to 15% by 47.5%. The welfare cost is concave. As OOT ownership grows, there are fewer local home owners and more renters. This tilts the welfare function more negative and erodes the political support for pro-OOT housing policies.

5.4.2 Persistence OOT Demand

In the baseline calibration, each OOT state has an expected duration of 40 years. We study a faster mean-reverting OOT process with persistence of $\pi = 0.75$ (16 year duration) and a slower one with $\pi = .95$ (80 years). Higher persistence strengthens the welfare effect. But even when we reduce expected length of each OOT regime from 40 to 16 years, the aggregate welfare cost only moderates from -0.74% to -0.70%. The lower-duration model features a much smaller construction response to an OOT inflow. The smaller housing supply response results in a fall in home ownership among the locals (and a smaller increase in the ownership rate overall). As persistence rises, changes in house prices become more similar to changes in rent because the mean reversion in rents is slower.

5.4.3 OOT Buyers Rent Out Property

The inability or unwillingness of OOT buyers to rent out their properties to locals is key for the welfare cost. In a model where OOT buyers rent out all of their real estate to local residents (Panel D), the welfare loss is close to zero. Rent and price increases are de minimis. Renters, who previously rented from local home owners, now rent from OOT buyers. As long as rents do not change much, renters are indifferent. The local real estate investors who have been displaced by OOT buyers now invest in the risk free asset. Because the housing risk premium is small, the risk free asset has a very similar rate of return to real estate, so local investors are also indifferent. Thus prices, quantities, and welfare are nearly unaffected; only the investment portfolio of some local investors changes. This result highlights that the main welfare cost of OOT investors is their decision to keep their property vacant. Leaving aside monitoring and compliance issues, "ghost" or "empty home" taxes like the ones introduced in Sydney and Vancouver, directly addresses the vacancy issue and could eliminate the welfare cost.

5.4.4 No Patience Heterogeneity

When all agents have the same subjective time discount factor ($\beta^H = \beta^L$=0.91, chosen to match the average wealth to average income ratio of the baseline model and the data), the model generates an aggregate welfare cost of OOT purchases of 0.37%, half as large as the 0.74% cost of the baseline. The effect on prices and construction is not very different from the baseline. The welfare cost for renters is substantially lower in the homogenous-β model. This model substantially understates wealth inequality. It generates too few poor renters and therefore too few losers from the OOT influx. The exercise underscores the importance of generating a realistic wealth distribution for assessing the welfare cost of increasing OOT home investment.

5.4.5 LTV limit

Panel F shows the result of an exercise that tightens the LTV limit for purchase mortgages (i.e., on primary homes) from the baseline value of 90% to 80% and an exercise that relaxes it to 95%. The welfare effect is somewhat less negative when LTV constraints are looser, but the effect is quantitatively small. In the case of tighter constraints, although the welfare loss is larger than the baseline model, renters lose by less than the baseline and owners gain by more. This is because with tighter constraints, the home ownership rate is lower. The ownership rate in the OOT^L state falls from 64.0% to 60.8% to 57.2% as the maximum LTV changes from 95% to 90% to 80%. This affects the composition of winners and losers from the OOT inflow.

5.4.6 Housing supply elasticity

Our baseline model generates a housing supply elasticity of 1.0 ($\overline{H}$=3), matching that in the average U.S. MSA. Panel G studies a case with much lower supply elasticity of 0.63 ($\overline{H}$=1) and with higher supply elasticity of 1.25 ($\overline{H}$=1,000). Miami's supply elasticity is close to 0.63, while Baltimore's is close to 1.25 (Saiz, 2010). The welfare effects are substantially larger in absolute value when housing supply is more inelastic. The reason is that rents increase by more in the inelastic economy, resulting in much larger welfare costs for renters. On impact, rents rise 19.59% vs. 18.77% in the benchmark. The long-run rent differences with the baseline model are larger (+14.79% vs. +13.58%). Over time, supply rises to accommodate the OOT purchases but it rises by far less in the inelastic housing model. The second experiment, which increases the amount of available residential land, thereby increasing the elasticity of housing supply, looks similar to the baseline suggesting that the baseline model is fairly close to the unconstrained case. This exercise shows that there can be substantial regional variation in the cost of living and welfare effects from OOT purchases, with much larger effects in supply-constrained cities like San Francisco, Miami, or New York.

5.4.7 Housing demand elasticity

Panel G studies two cases with a less elastic housing demand than the baseline, with intratemporal elasticities of substitution between housing and non-housing consumption of 0.33 and 0.50, and one case with an elasticity of 1, corresponding to a Cobb-Douglas aggregator. The baseline model has an elasticity of 0.67. The literature has found a wide range of estimates for this parameter.[15] The welfare costs of OOT inflows are substantially higher when housing demand is more inelastic. The aggregate welfare cost is 1.08% for an elasticity of 0.33 and 0.92% when the elasticity is 0.50. Because households are less willing to adjust quantities, rent and house prices must do much more of the adjustment. The OOT inflow triggers a rent (house price) increase of 30.61% (15.24%) when the elasticity is 0.33. Both are much larger than in the baseline calibration. The opposite is true with Cobb-Douglas preferences. Rents and house prices increase by less upon an OOT inflow as households are more willing to substitute between housing and non-housing consumption. This limits the welfare cost to renters and the welfare gain to owners. The net effect is a lower aggregate welfare cost.

5.4.8 No minimum housing size

The baseline model features a minimum housing size equal to one-third of the average house size. Panel I reports a version of the model without this constraint. The aggregate welfare cost drops by nearly half from 0.74% to 0.38%. Price and even rent changes are not affected much, in contrast, because the wealth-weighted average owner and renter

[15] For example, Hanushek and Quigley (1980) report values of 0.45-0.65 for a range of cities, Stokey (2009) uses 0.5 based on CEX data for 27 cities, Li, Liu, Yang, and Yao (2016) estimate 0.48, and the CBO uses 0.4. Flavin and Nakagawa (2008) uses the PSID but a different measure of housing and estimate a value of 0.13.

are not much affected by the minimum house size constraint. The large change in welfare is mostly due to (poor) renters. Renters who are unconstrained by a minimum house size requirement can reduce housing consumption in the face of an increase in rents, thereby limiting the adverse welfare effect. A minimum house size requirement hurts the poorest households by eliminating this margin of adjustment.

5.4.9 Profit redistribution

The baseline model assumes that firm profits leave the city. Panel J studies an economy where half of firm profits are redistributed lump-sum to the local residents. Profit redistribution has little effect on the aggregate welfare effect. The profit redistribution cushions the welfare cost from the OOT inflow to renters and increases the welfare benefit to owners. This positive effect on welfare is modest because aggregate profits fall when OOT investors enter. Profit increases in the construction sector are more than offset by profit reductions in the tradable goods sector. Profit redistribution also substantially lowers the baseline home ownership rate, and therefore the composition of winners and losers from the OOT inflow. The resulting welfare effect is only slightly less negative than without profit redistribution.

5.4.10 Owners Only

Finally, we ask whether there is a set of parameters for which the welfare effects of OOT purchases are a net positive to the city. Panel K presents an example of such an economy with a small 0.04% benefit. In other words, the negative welfare effect is not a generic outcome, just the most likely outcome under a realistically calibrated model. The model requires several strong assumptions that bring the model closer to a representative agent economy: (i) all local residents must be home owners, (ii) no patience heterogeneity, (iii) no minimum housing size, (iv) no bequest heterogeneity, and (v) all bequests are received at age 21. The last assumption is necessary if renting is not an option. Otherwise, some households would enter the model with zero wealth at age 21 and could not afford to own a house, violating assumption (i). With these changes, the capital gain effect dominates and the welfare effect is positive. This case illustrates the importance of allowing for tenure choice and generating realistic wealth inequality, including through bequests.

6 Policy Responses

Having established a net welfare loss from OOT purchases, a natural question is whether policy can offset these welfare costs. So far, we have considered two policy responses. The first one was to force OOT buyers to rent out their property. We showed in Panel D of Table 2 that this policy eliminates the welfare cost. Second, we considered abolishing a minimum housing size requirement. That policy was also able to lower the welfare costs substantially (Panel I of Table 2). In this section we consider three additional policy responses, motivated by the empirical evidence discussed in section 2. First, we study a policy that levies extra taxes on OOT buyers. To render the exercise more realistic and meaningful, we introduce two new model ingredients. We make OOT purchases sensitive to the price of real estate. And we introduce a public good which provides utility to the local residents. The extra revenue raised by the tax on OOT investors increases public good provision. In the second policy experiment we also tax OOT investors but lower property taxes on local residents in such a way that leaves total tax revenue unchanged. This experiment does not require us to take a stance on the utility over public goods. The third experiment considers a change in housing supply regulation in response to the OOT inflow. All three policies can fully offset the welfare costs from OOT demand, but have different distributional implications.

6.1 Nature of OOT Tax

We model the OOT tax as a property tax surcharge, τ^{OOT}. Such a tax on OOT buyers was introduced in Vancouver and Sydney recently. We calculate the present-value equivalent of the tax by taking into account the time value of money as well as the transition probabilities between the OOT regimes. For example, a tax rate of $\tau^{OOT} = 3.46\%$ per four years, or 0.86% per year, is equivalent to a one-time 15% transaction tax. Adding the surcharge to the property tax rate of 4% charged to locals (1% per year), results in a OOT tax rate of 7.46%. We explore a wide range of values from 4% (baseline, no surcharge) to 14% per four years.

6.2 Price-Elastic OOT Purchases

The OOT purchases continue to follow a two-state Markov chain. As before, $OOT^L = 0$ in the low state. OOT purchases in the high state now depend on the equilibrium house price:

$$\log\left(OOT_t^H\right) = a - b\log\left(P_t(1+\tau^{OOT})\right) \tag{9}$$

When the price elasticity of OOT demand $b = 0$, we recover the baseline model. We consider two values for b. The first one sets $b = 6.4$ to match the observed change in foreign home purchases in Vancouver in response to the introduction of the transaction tax, which effectively increased prices on foreigners by 15%. The share of OOT purchases in the five weeks prior to the OOT tax was 10%.[16] After the tax was introduced the OOT share fell to 4.1% of purchases.[17] The implied price elasticity is $(\log(0.041) - \log(0.10))/(\log(1.15) - \log(1)) = 6.4$. We set a to deliver the same unconditional OOT demand in the model with $b = 6.4$ as in the baseline price-inelastic model. Suher (2016) estimates a much lower price elasticity of 0.6 based on a 2013 experiment in New York City when a property tax abatement was phased out for non-residents. We use this much lower price elasticity estimate as our second value for $b = 0.6$.

6.3 Local Public Goods

We change the utility function to give local residents additional utility over a local public good G. This public good is funded with property tax revenue. The additional property tax revenue coming from the surcharge on OOT buyers increases the supply of the local public good. The period utility function we use is:

$$F(G_t)U(c_t, h_t, n_t) = \exp((1-\gamma)\theta G_t) \times \frac{\mathcal{C}(c_t, h_t, l_t)^{1-\gamma}}{1-\gamma} \tag{10}$$

The parameter $\theta \geq 0$ modulates the importance of the public good in the utility function. It captures the government's efficacy at converting a given amount of tax revenue into valuable public goods. The case of $\theta = 0$ recovers our baseline model without utility over public goods (wasteful spending of property tax revenue). The public finance literature provides little guidance on how to set θ. We explore a range of values for θ between 0 and 2. Values of θ below 1 imply that the average household is unwilling to forgo private consumption for more public goods. Values of θ above 2 imply that households

[16] Official statistics indicate a 10% share for the period from June 10 until July 16, 2016. This period excludes the two weeks prior to the announcement of the tax and another month prior to the introduction of the tax, during which anticipation buying may have artificially inflated the foreign purchase share. Several articles indicate that a foreign purchase share of 10-13% was common for metro Vancouver in the months before June 10, 2016. E.g., http://www.rew.ca/news/foreign-purchases-of-vancouver-homes-rise-slightly-as-market-adjusts-to-tax-1.3412049.

[17] The 4.1% share is for January 2017. Because of the initial bunching of OOT purchases prior to August 2016, there was a sharp drop-off in the immediate aftermath of the introduction of the tax. The foreign share was only 1.8% in September 2016 and gradually started to increase to 3-4% in the October 2016-January 2017 period. Foreign purchases fell from 4.1% to 2.5% by mid-2017 before starting to rise again, but the post-January 2017 changes are likely due to reasons other than the transaction tax.

want to forgo so much private consumption that optimal government tax revenue would end up above 50% of private consumption. In the model with $\tau^{OOT} = 3.46\%$, property tax revenue is about 5% of private consumption. This level of tax revenue is optimal for θ between 1.10 and 1.15, and close to the observed share of local tax revenue in Vancouver's GDP.

6.4 Welfare Effects from OOT Tax

Table 3 shows the results for the main tax experiment for various tax rates on OOT buyers τ^{OOT}, presented in the rows of the table, and different values for the utility benefit from public goods θ, presented in the columns. The third column reports the one-time "transaction" tax that is equivalent in present value terms to the per-period tax rate in the first column. The table reports the aggregate welfare effect $\mathcal{W}$ of OOT purchases, expressed in percent consumption equivalent. The (0,0) entry in the table is similar to the baseline model, except for the non-zero price elasticity of OOT purchases.

Table 3: Welfare Effects From Vancouver Tax Policy

OOT tax			Panel A: Price elasticity $b = 6.4$							Panel B: Price elasticity $b = 0.6$						
τ^{OOT}		PV-equivalent	Public Good Utility Parameter θ							Public Good Utility Parameter θ						
4yr	1yr	transaction tax	0	0.5	0.75	1	1.25	1.5	2	0	0.5	0.75	1	1.25	1.5	2
0%	0%	0%	-0.66	-0.57	-0.51	-0.46	-0.40	-0.35	-0.23	-0.72	-0.63	-0.57	-0.51	-0.46	-0.40	-0.28
1%	0.25%	4.4%	-0.66	-0.53	-0.46	-0.39	-0.33	-0.26	-0.13	-0.72	-0.59	-0.52	-0.45	-0.38	-0.30	-0.16
2%	0.5%	8.7%	-0.65	-0.50	-0.42	-0.34	-0.26	-0.18	-0.03	-0.72	-0.56	-0.47	-0.39	-0.30	-0.21	-0.04
3.46%	0.86%	15.0%	-0.63	-0.45	-0.36	**-0.27**	**-0.18**	-0.09	0.09	-0.72	-0.51	-0.40	-0.29	-0.19	-0.08	0.12
5%	1.25%	26.1%	-0.60	-0.40	-0.29	-0.19	-0.09	0.01	0.21	-0.72	-0.46	-0.33	-0.20	-0.07	-0.05	0.29
7.5%	2.5%	34.8%	-0.54	-0.30	-0.18	-0.06	0.06	0.18	0.40	-0.71	-0.38	-0.22	-0.06	0.10	0.25	0.57
10%	3.0%	52.2%	-0.43	-0.17	-0.04	0.09	0.22	0.34	0.59	-0.70	-0.30	-0.11	0.08	0.27	0.46	0.85

Notes: The table reports the net aggregate welfare effect $\mathcal{W}$ of an increase in OOT purchases from the low state to the high state, aggregated across all the agents in the economy. The numbers represent a consumption equivalent variation, expressed in percent. Each entry is for a different combination of the OOT tax rate τ^{OOT} (per 4 years), listed in the first column, and the utility parameter for local public goods θ, listed in the top row. Note that residents pay a property tax rate of τ^P while OOT investors pay a property tax rate of $\tau^P + \tau^{OOT}$. The second column expresses the tax rate surcharge into an annual number while the third column reports the present-value equivalent one-time transaction tax rate. Panel A assumes that OOT housing demand in the high state has a price elasticity of 6.4, while panel B assumes a price elasticity of OOT demand in the high state of 0.6.

The main result is that introducing a tax on OOT buyers is welfare increasing. As the OOT tax surcharge goes from 0% to 10% per four-year period, the aggregate welfare cost of OOT buyers to local residents falls for every value of θ. When the extra tax revenue the tax brings in is wasted ($\theta = 0$), the welfare cost of OOT buyers falls from 0.66% to 0.43% as taxes on OOT buyers are raised. The welfare effects are concave in the OOT tax rate. When taxes become very high, the share of OOT purchases becomes very small, especially when $b = 6.4$. By halting (most) of the OOT inflow, high taxes mitigate the adverse consequences of the OOT purchases. This leaves more housing for residents and reduces the rent. Average renter welfare increases from -2.20% to (not reported in the table). In contrast, owners' welfare falls from +0.53% to +0.31% as capital gains become smaller when fewer OOT buyers enter.

Reading across the columns, the greater the utility value locals derive from public goods, the higher welfare for a given tax rate. As θ increases from 0 to 2, a 15% transaction tax (row denoted 7.46%) turns a -0.63% welfare cost into a +0.09% welfare gain. The locus of (τ^{OOT}, θ) combinations for which OOT purchases are welfare neutral is downward sloping. The more locals value public goods, the more they value the additional tax revenue raised by taxing OOT investors, and the more steeply welfare increases in the tax rate. At the empirically most relevant range $\theta \in [1.1, 1.25]$, a 15% transaction tax reduces the aggregate welfare loss from about -0.43% without the tax to -0.23% with the tax (interpolating the number in bold in the table). This suggests the Vancouver policy was quite effective at reducing the welfare cost of the OOT inflow.

Panel B of Table 3 explores how sensitive the results are to the price elasticity of OOT demand. For $b = 0.6$, OOT demand does not decline as much when the tax rate increases. Put differently, it takes a much larger tax to dissuade OOT investors from buying. This economy is much closer to our benchmark model with completely price inelastic purchases. Indeed, the aggregate welfare effect of -0.72% in the (0,0) entry in Panel B is very similar to the -0.74% welfare cost presented for our benchmark model. For $\theta = 0$, welfare increases from -0.72% to -0.70% as the OOT tax increases from 0% to 10% per four years. For $\theta \in [1.1, 1.25]$, a 15% transaction tax reduces the aggregate welfare loss by two-thirds, from about -0.49% without the tax to about -0.24% with the tax. With less elastic OOT purchases, owners gain more from the tax increase than renters since capital gains remain strong and rents do not fall much. We conclude that the aggregate welfare results are rather insensitive to the price elasticity of OOT demand, but the distribution of the welfare gains from the tax differs.

6.5 Revenue-neutral Property Taxation

In a second tax experiment, we levy a property tax surcharge on OOT buyers and simultaneously reduce property taxes on local residents so as to keep total tax revenue constant before and after the OOT inflow. This exercise does not require us to take a stance on the utility value of public goods θ. The first row of Table 4 reports an economy where the property tax rate is 4% in both states of the world and residents and non-residents are taxed equally. This is the baseline economy, except for the fact that OOT housing demand is now price-elastic. In Panel A, we use the high price elasticity estimate of $b = 6.4$, while in Panel B we use the low estimate $b = 0.6$. Property tax revenue is higher in the high

OOT state because of a larger housing stock and higher house prices.

Row 2 shows that we can reduce the property tax on locals to 3.6% while keeping total tax revenue constant before and after the OOT inflow. From row 3 down, the property tax on OOT investors is raised while taxes on locals are lowered to keep tax revenue constant. Naturally, when OOT demand is more price elastic, OOT demand in the high state falls more as taxes rise. As a result, taxes on residents cannot be lowered as much as when OOT demand is less price elastic. Lower property taxes are a boon for home owners and amplify the gains from the OOT inflow observed in the baseline economy. The positive welfare effect for owners is larger when OOT demand is more inelastic (in Panel B) because the property tax reduction is larger and the capital gains are larger. High OOT taxes are also good news for local renters because they reduce the size of the OOT capital inflow, thereby increasing the available housing for locals. The positive effects on local renters are smaller when OOT demand is more inelastic (in Panel B), since there is less of a reduction in the OOT demand in response to the higher taxes.

Higher OOT taxes increase resident welfare, benefitting both renters and owners. A zero welfare effect can be obtained for a high enough OOT tax. That tax rate is 11.5%, corresponding to an annual property tax surcharge of 1.865%, in Panel B. The welfare-neutral tax rate is slightly lower in Panel A. When OOT demand is more price elastic (Panel A), renters' welfare increases sharply as taxes are increased. When OOT demand is more price inelastic (Panel B), owners' welfare increases faster with OOT tax rates.

A 15% transaction tax (the row 7.46%) offsets two-thirds of the welfare losses from the OOT inflow for both values of b, similar to the first tax experiment.

Table 4: Welfare Effects From Alternative Tax Policy

	Panel A: Price elasticity $b = 6.4$				Panel B: Price elasticity $b = 0.6$			
$\tau^P + \tau^{OOT}$	τ^P	$\mathcal{W}$	$\mathcal{W}_r$	$\mathcal{W}_o$	τ^P	$\mathcal{W}$	$\mathcal{W}_r$	$\mathcal{W}_o$
4%	4%	-0.66	-2.20	0.53	4%	-0.72	-2.40	0.56
4%	3.60%	-0.45	-2.12	0.85	3.60%	-0.52	-2.34	0.91
5%	3.46%	-0.34	-1.93	0.92	3.44%	-0.42	-2.27	1.05
6%	3.40%	-0.29	-1.83	0.94	3.30%	-0.33	-2.23	1.17
7.46%	3.35%	-0.23	-1.70	0.94	3.16%	-0.24	-2.17	1.30
9.%	3.25%	-0.11	-1.53	1.01	3.02%	-0.15	-2.12	1.42
11.5%	3.15%	0.06	-1.26	1.11	2.77%	0.00	-2.03	1.64
14%	3.08%	0.19	-1.03	1.18	2.60%	0.17	-1.94	1.87

Notes: The table reports the effects of charging property taxes on OOT investors, listed in the first column. Except for the first row, the property taxes charged on residents, τ^P, listed in columns 2 and 6 are such that property tax revenue is unchanged between the low OOT state and the high OOT state. In the first row, we report a baseline case where residents and OOT investors face the same tax rate. That case has larger tax revenue in the high OOT state compared to the low OOT state. The price elasticity of OOT demand in the high state is $b = 6.4$ in panel A and $b = 0.6$ in panel B. The table reports the aggregate welfare effect from a change from the low OOT to the high OOT state, $\mathcal{W}$, the average welfare effect for renters $\mathcal{W}_r$, and the average welfare effect for owners $\mathcal{W}_o$. The welfare effects are expressed as percentages of consumption equivalent units.

6.6 Relaxing Housing Supply Regulation

A different policy response to an OOT influx is to facilitate a stronger housing supply response. An increase in $\overline{H}$ can be interpreted as a relaxation of the regulatory barriers to construction (more and faster permitting, more permissive zoning, etc.). To capture such a policy response, we consider a model where $\overline{H}$ is 3 in the low OOT state and 25 in the high OOT state. The model is otherwise identical to the benchmark model. This change in $\overline{H}$ corresponds to an increase in the elasticity of housing supply from 1.00 in the low state to 1.22 in the high state. Put differently, the increase in OOT demand is met by an increase in the housing stock of 7.9%, or 80% of the OOT demand. In the baseline model, the increase in the housing stock is 2.8%, accommodating only 28% of the OOT demand. This supply response is chosen to generate an aggregate welfare effect from the OOT increase of exactly 0.0%. The supply response that eliminates the welfare cost of OOT investors is akin to changing the city's housing supply elasticity from that of Tampa (1.0) into that of Baltimore (1.23) or Tacoma (1.21).

While this policy response can generate the same welfare-neutral outcome in the aggregate as for the tax policies, it has different distributional implications. Most of the welfare gains accrue to renters. Renters only suffer a 0.40% welfare loss in this economy, compared to a 2.43% loss in the benchmark model. Owners' gains shrink to 0.29%, compared to 0.56% in the benchmark. In the tax policy exercises with price inelastic OOT demand, the welfare benefits instead disproportionately accrued to owners.

7 Conclusions

Recent years have seen a surge in out-of-town (OOT) home buyers in many gateway cities all over the world. This paper develops an equilibrium model that allows for a quantitative welfare analysis of such OOT purchases. The model not only matches patterns of wealth accumulation and home ownership over the life-cycle, but also results in realistic house prices, rents, and wages for the city. When OOT buyers represent 10% of housing demand and do not rent out their properties, they impose a net welfare cost on the city's residents of 0.74%. A tax on OOT investors, calibrated to the 2016 Vancouver tax, offsets two-thirds of this welfare cost. Relaxing housing supply regulations also provides an effective means to offset the welfare losses.

In future work, we plan to study the aggregate and distributional implications of housing affordability policies, such as zoning and rent control, that aim to address the high cost of housing in the major cities of the world. This work will extend the model to multiple neighborhoods within the city, enabling us to make contact with the data on inequality within and across neighborhoods, and allowing for a richer analysis of place-based policies.

References

AGARWAL, S., T. F. SING, AND L. WANG (2018): "Information Asymmetries and Learning in Commercial Real estate Markets," Working Paper Georgetown University.

AMERIKS, J., A. CAPLIN, S. LAUFER, AND S. VAN NIEUWERBURGH (2011): "The Joy of Giving or Assisted Living? Using Strategic Surveys to Separate Public Care Aversion from Bequest Motives," *Journal of Finance*, 66, 519–561.

BADARINZA, C. AND T. RAMADORAI (2018): "Home Away From Home? Foreign Demand and London House Prices," *Journal of Financial Economics*, forthcoming.

BADARINZA, C., T. RAMADORAI, AND C. SHIMIZU (2018): "Nationality bias at home and abroad: Theory and evidence from commercial real estate," Working Paper.

BAYER, P., C. GEISSLER, K. MAGNUM, AND J. W. ROBERTS (2011): "Speculators and Middlemen: The Strategy and Performance of Investors in the Housing Market," Working Paper Duke University.

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 (2017): "House prices and consumer spending," *Review of Economic Studies*, 85, 1502–1542.

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

CHINCO, A. AND C. MAYER (2016): "Misinformed Speculators and Mispricing in the Housing Market," *Review of Financial Studies*, 29, 486–522.

COCCO, J. AND P. LOPES (2017): "Reverse Mortgage Design," *Working Paper London Business School and London School of Economics*.

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

CVIJANOVIC, D. AND C. SPAENJERS (2018): "Well Always Have Paris: Out-of-Country Buyers in the Housing Market," Working Paper.

DAVIS, M. A. AND S. VAN NIEUWERBURGH (2015): *Handbook of Regional and Urban Economics*, North Holland, chap. Housing, Finance, and the Macroeconomy, Chapter 12.

DEFUSCO, A., C. NATHANSON, AND E. ZWICK (2017): "Speculative Dynamics of Prices and Volume," Working Paper.

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.

FLAVIN, M. AND S. NAKAGAWA (2008): "A model of housing in the presence of adjustment costs: A structural interpretation of habit persistence," *American Economic Review*, 98, 474–495.

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.

HANUSHEK, E. A. AND J. M. QUIGLEY (1980): "What Is the Price Elasticity of Housing Demand?" *The Review of Economics and Statistics*, 62, 449–454.

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 (2017): "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.

LI, W., H. LIU, F. YANG, AND R. YAO (2016): "Housing over Time and over the Life Cycle: A Structural Estimation," *International Economic Review*, 1237–1260.

ORTALO-MAGNÉ, F. AND A. PRAT (2016): "Spatial Asset Pricing: A First Step," *Economica*, 83, 130–171.

ROSSI-HANSBERG, E., P. D. SARTE, AND R. OWENS (2010): "Housing Externalities," *Journal of Political Economy*, 118, 485–535.

RUF, T. AND M. D. LEVI (2011): "The Law of One Price in Unfamiliar Places: The Case of International Real Estate," Working Paper UBC Sauder School.

SÁ, F. (2017): "The Effect of Foreign Investors on Local Housing Markets: Evidence from the UK," Working Paper.

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

STOKEY, N. L. (2009): "Moving Costs, Nondurable Consumption and Portfolio Choice," *Journal of Economic Theory,* forthcoming.

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.

SUHER, M. (2016): "Is Anybody Home? The Impact and Taxation of Non-Resident Buyers," Working Paper Furman Center New York University.

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.

WHEATON, W. (2017): "The Determinants of Housing Price Volatility; 2001-2016," Working Paper MIT.

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

A Analytical solution for housing and labor supply choices

We will solve only the worker's problem here. A retiree's problem is analogous, but simpler because there is one fewer choice as $n_t = 0$.

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, ξ_t be the Lagrange multiplier on the labor constraint, and μ_t be the Lagrange multiplier on the housing constraint. The numerical strategy is to choose c_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 c_t.

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

$$\begin{aligned}
&\mathcal{C}_t^{1-\gamma-\eta}(1-\alpha_N)(1-\alpha_H)((1-\alpha_H)c_t^{\epsilon}+\alpha_H h_t^{\epsilon})^{\frac{\eta-\epsilon}{\epsilon}}c_t^{\epsilon-1} = \lambda_t \\
&\mathcal{C}_t^{1-\gamma-\eta}(1-\alpha_N)\alpha_H((1-\alpha_H)c_t^{\epsilon}+\alpha_H h_t^{\epsilon})^{\frac{\eta-\epsilon}{\epsilon}}h_t^{\epsilon-1} = \lambda_t R_t \\
&-\mathcal{C}_t^{1-\gamma-\eta}\alpha_N\chi(1-n_t)^{\chi\eta-1} = \lambda_t G^a W_t(1-\tau^{SS}) \\
&\lambda_t = Q\beta E_t[\tfrac{\partial V_{t+1}}{\partial x_{t+1}}]
\end{aligned} \tag{11}$$

By rearranging, it can be shown that conditional on choosing non-durable consumption c_t,

$$h_t = \left(\frac{\alpha_H}{1-\alpha_H}\right)^{\frac{1}{1-\epsilon}} R_t^{\frac{-1}{1-\epsilon}} c_t \tag{12}$$

and

$$n_t = 1 - c_t^{\frac{1-\eta}{1-\chi\eta}} \left(\frac{1}{(1-\tau)W_t G_a z_t}\right)^{\frac{1}{1-\chi\eta}} \left(\frac{\alpha_N \chi}{(1-\alpha_N)(1-\alpha_H)}\right)^{\frac{1}{1-\chi\eta}} \left(\frac{(1-\alpha_H)R_t}{\alpha_H}\right)^{\frac{\alpha_H \eta}{1-\chi\eta}}. \tag{13}$$

Case 2: $\nu_t > 0$, $\xi_t = 0$, and $\mu_t = 0$. In this case the borrowing constraint binds and $b_{t+1} = 0$, but the labor constraint does not. The first order conditions in the first three lines of equation (11) are still correct. It is still the case that conditional on choosing c_t, h_t and n_t are described by the equations above. By plugging these into the budget constraint, it is also possible to explicitly solve for non-durable consumption.

Case 3: $\nu_t = 0$ and $\xi_t > 0$, and $\mu_t = 0$. In this case the borrowing constraint does not bind, but the labor constraint does, implying $n_t = \underline{n}$. The first order conditions in the first, second, and fourth lines of equation (11) are still correct. As in case 1, conditional each choice c_t, h_t is described by the equation above.

Case 4: $\nu_t > 0$ and $\xi_t > 0$, and $\mu_t = 0$. In this case both constraints bind, implying $n_t = \underline{n}$ and $b_{t+1} = 0$. The first order conditions in the first and second lines of equation (11) are still correct, as is the equation for h_t as a function of c_t. By plugging this into the budget constraint, it is also possible to explicitly solve for non-durable consumption.

Case 5: $\nu_t = 0$, $\xi_t = 0$, and $\mu_t > 0$. In this case the minimum housing constraint binds. We set $h_t = \underline{h}$ but the first order conditions for consumption and labor are unaffected, so equation (13) still holds. If this new housing choice violates the borrowing constraint, then conditional on each choice of c_t, we increase n_t until the borrowing constraint is satisfied. If the borrowing constraint is still not satisfied, we assign a large negative utility to this choice of c_t.

Case 6: $\nu_t > 0$, $\xi_t = 0$, and $\mu_t > 0$. In this case, both the minimum housing constraint and the borrowing constraint bind. Conditional on each choice of c_t, we set $h_t = \underline{h}$ and we increase n_t until the borrowing constraint is satisfied. If the borrowing constraint is still not satisfied, we assign a large negative utility to this choice of c_t.

Case 7: $\nu_t = 0$ *and* $\xi_t > 0$, *and* $\mu_t > 0$. In this case the labor constraint and housing constraint both bind. Conditional on each choice of c_t, we set $n_t = \underline{n}$ and $h_t = \underline{h}$. If this new housing choice violates the borrowing constraint, then we increase n_t until the borrowing constraint is satisfied. If the borrowing constraint is still not satisfied, we assign a large negative utility to this choice of c_t.

Case 8: $\nu_t > 0$ *and* $\xi_t > 0$, *and* $\mu_t > 0$. In this case all constraints bind. Identically to Case 7, we set $n_t = \underline{n}$ and $h_t = \underline{h}$. By using the budget constraint, it is possible to explicitly solve for non-durable consumption. If the necessary consumption is negative, we assign a large negative utility to this individual state.

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, ξ_t be the Lagrange multiplier on the labor constraint, and μ_t be the Lagrange multiplier on the housing constraint. We also extend the baseline model in the text by allowing rental property to depreciate at an additional rate δ^I compared to owner-occupied property. The numerical strategy is to choose c_t and $\widehat{h}_t$ 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 c_t and $\widehat{h}_t$.

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

$$\begin{aligned}
&\mathcal{C}_t^{1-\gamma-\eta}(1-\alpha_N)(1-\alpha_H)((1-\alpha_H)c_t^{\epsilon}+\alpha_H h_t^{\epsilon})^{\frac{\eta-\epsilon}{\epsilon}}c_t^{\epsilon-1} = \lambda_t \\
&\mathcal{C}_t^{1-\gamma-\eta}(1-\alpha_N)\alpha_H((1-\alpha_H)c_t^{\epsilon}+\alpha_H h_t^{\epsilon})^{\frac{\eta-\epsilon}{\epsilon}}h_t^{\epsilon-1} + \beta E_t[\tfrac{\partial V_{t+1}}{\partial x_{t+1}}P_{t+1}(1-\delta-\tau^P)] = \lambda_t P_t \\
&-\mathcal{C}_t^{1-\gamma-\eta}\alpha_N\chi(1-n_t)^{\chi\eta-1} = \lambda_t G^a W_t(1-\tau^{SS}) \\
&\lambda_t = Q\beta E_t[\tfrac{\partial V_{t+1}}{\partial x_{t+1}}] \\
&\beta E_t[\tfrac{\partial V_{t+1}}{\partial x_{t+1}}P_{t+1}(1-\delta-\tau^P)] = \lambda_t(P_t - R_t)
\end{aligned} \tag{14}$$

By rearranging, it can be shown that the equations for residential housing choice and hours are identical to equations (12) and (13) above.

Case 2: $\nu_t > 0$ *and* $\xi_t = 0$, *and* $\mu_t = 0$. In this case the borrowing constraint binds implying $x_t - P_t\widehat{h}_t - P_t h_t = -P_t(\theta_{res}h_t + \theta_{in}\widehat{h}_t)$, but the labor constraint does not bind. Eliminating b_{t+1} from the budget constraint, we can solve for the residential housing choice as a function of $(c_t, \widehat{h}_t)$:

$$h_t = (x_t - (1-\theta_{inv})P_t\widehat{h}_t)/((1-\theta_{res})P_t). \tag{15}$$

The hours choice remains as in equation (13).

Case 3: $\nu_t = 0$ *and* $\xi_t > 0$, *and* $\mu_t = 0$. In this case the borrowing constraint does not bind, but the labor constraint does, implying $n_t = \underline{n}$. All but the third line of equation (14) are still correct. Conditional on choosing $(c_t, \widehat{h}_t)$, h_t is identical to Case 1.

Case 4: $\nu_t > 0$ *and* $\xi_t > 0$, *and* $\mu_t = 0$. In this case both constraints bind, implying $n_t = \underline{n}$ and $x_t - P_t\widehat{h}_t - P_t h_t = -P_t(\theta_{res}h_t + \theta_{in}\widehat{h}_t)$. Residential housing is identical to equation (15) in Case 2.

Case 5: $\nu_t = 0$ *and* $\xi_t = 0$, *and* $\mu_t > 0$. In this case the minimum housing constraint binds. We set $h_t = \underline{h}$ but the first order conditions for consumption and labor are unaffected, so equation 13 still holds. If this new housing choice violates the borrowing constraint, then conditional on each choice of $(c_t, \widehat{h}_t)$, we increase n_t until the borrowing constraint is satisfied. If the borrowing constraint is still not satisfied, we assign a large negative utility to this choice of $(c_t, \widehat{h}_t)$.

Case 6: $\nu_t > 0$ *and* $\xi_t = 0$, *and* $\mu_t > 0$. In this case, both the minimum housing constraint and the borrowing constraint bind. Conditional on each choice of $(c_t, \widehat{h}_t)$, we set $h_t = \underline{h}$ and we increase n_t until the borrowing constraint is satisfied. If the borrowing constraint is still not satisfied, we assign a large negative utility to this choice of $(c_t, \widehat{h}_t)$.

Case 7: $\nu_t = 0$ and $\xi_t > 0$, and $\mu_t > 0$. In this case the labor constraint and housing constraint both bind. Conditional on each choice of $(c_t, \widehat{h}_t)$, we set $n_t = \underline{n}$ and $h_t = \underline{h}$. If this new housing choice violates the borrowing constraint, then we increase n_t until the borrowing constraint is satisfied. If the borrowing constraint is still not satisfied, we assign a large negative utility to this choice of $(c_t, \widehat{h}_t)$.

Case 8: $\nu_t > 0$ and $\xi_t > 0$, and $\mu_t > 0$. In this case all constraints bind. Identically to Case 7, we set $n_t = \underline{n}$ and $h_t = \underline{h}$. By using the budget constraint, it is possible to explicitly solve for non-durable consumption. If the necessary consumption is negative, we assign a large negative utility to this individual state.

B Labor Income Calibration

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

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

where G^a is a function of age and $z_{a,i}$ is the idiosyncratic component of productivity.

The average idiosyncratic productivity, conditional on age, is one: $E[z_{a,i}|a] = 1$. We let $z_{a,i}$ depend on age a because, in the data, the cross-sectional dispersion of income grows with age. For a given age a, $z_{a,i}$ takes on one of three values (low, middle, high), but these values differ by age. Idiosyncratic productivity $z_{a,i}$ is stochastic, and its transition dynamics are governed by a Markov process $\mathcal{P}(i, i'; \beta)$. Transition probabilities depends on i but not on a. Furthermore, we allow the transition probability matrix to differ for households with low saving motive ($\beta = \beta_L$) and high saving motive ($\beta = \beta_H$), as detailed below. This is done in order to induce a sufficiently positive correlation between income and wealth in the model.

We define earnings as labor income plus business income in the Survey of Consumer Finances (SCF). We use 10 survey waves from 1983-2010. For the purpose of these calculations, we exclude households that report negative or zero income. Below, we explain how we calibrate the three components of the income process: (A) the age profile G^a; (B) the idiosyncratic income relative to mean $z_{a,i}$; (C) the transition probability density of z, $\mathcal{P}$.

(A) For each SCF wave, we compute average earnings in each 4-year age bucket (above age 21), and divide it by the average income of all households (above age 21). This gives us an average relative income at each age. We then average this relative age-income across all 10 SCF waves. This is G^a.

(B) To compute the idiosyncratic income $z_{a,i}$ of each group $i \in (low, middle, high)$ at a particular age a, relative to the average income of all households of that age we do the following:

Step 1: For each positive-earnings household, we compute whether it was in the bottom 25%, middle 50%, or top 25% of earners of the same age. Thus, each household is defined as a low, middle, or high earner (relative to same-aged households).

Step 2: For each 4-year age bucket, we compute average earnings of all low earners and define this as $z^{\$}_{a,j}$ where a is the 4-year age group and $j = low$. We do the same for $j = middle$ and $j = high$.

Step 3: We normalize each group's income by the average income in each age group, to get each group's relative income:

$$z_{a,j} = \frac{z^{\$}_{a,j}}{0.25z^{\$}_{a,low} + 0.5z^{\$}_{a,middle} + 0.25z^{\$}_{a,high}}.$$

Step 4: Steps 1-3 above are done separately for each wave of the SCF. We compute an equal-weighted average across all 10 waves to get an average relative income for each age and income group. This gives us three 11x1 vectors $z_{a,low}$, $z_{a,middle}$, and $z_{a,high}$ since there are 11 4-year age groups between ages 21 (entry into job market) and 65 (retirement). Note that the average Z across all households of a particular age group is always one: $E[z_{a,j}|a] = 1$.

Step 5: We regress each vector (low, middle, high), on a linear trend to get a linearly fitted value for each group's (low, middle, high) relative income at each age. The reason we perform Step 5, rather stopping at Step 4 is that the relative income at age 4 exhibits some small non-monotonicities that are likely caused by statistical noise (sampling and measurement error). Step 5 smoothes this out.

(C) Storesletten, Telmer, and Yaron (2006) estimate that the autocorrelation of the idiosyncratic component of labor income is 0.9. Thus, if theβ_L and β_H agents had the same transition probability matrix $\mathcal{P}$, we would set

$$\mathcal{P}(i,i') = \begin{bmatrix} p & 1-p & 0 \\ .5(1-p) & p & .5(1-p) \\ 0 & 1-p & p \end{bmatrix}$$

for each agent, with $p = 0.9$. This matrix implies that the low, middle, and high income groups are, respectively, 25%, 50%, and 25% of the population, and that the persistence of the idiosyncratic component is 0.9. However, setting this $\mathcal{P}$ for all agents results in a correlation between income and wealth that is too low compared to the data. For this reason, we adjust P in the following way. For β_H-agents, we raise the persistence of the high state, choosing $p_H > p$ in:

$$\mathcal{P}(i,i';\beta_H) = \begin{bmatrix} p & 1-p & 0 \\ .5(1-p) & p & .5(1-p) \\ 0 & 1-p_H & p_H \end{bmatrix}$$

We set $p_H = 0.95$ to approximately match the wealth share held by the top 25% of the income distribution. Note that this implies that among β_H agents, the low, middle, and high income groups are, respectively, 20%, 40%, and 40% of the population. In order to keep the unconditional fractions at 25%, 50%, and 25%, we must also change $\mathcal{P}$ for the β_L-agents. We set it to:

$$\mathcal{P}(i,i';\beta_L) = \begin{bmatrix} p & 1-p & 0 \\ .5(1-p) & p & .5(1-p) \\ 0 & 1-p_L & p_L \end{bmatrix}$$

where $p_L = 0.8667$ ensures that the unconditional fractions of low, middle, and high earners are 25%, 50%, and 25%. To sum up, we choose p, p_H, and p_L to match (i) the persistence of the idiosyncratic component of wages, (ii) the share of wealth held by the top 25% of income earners, and (iii) the unconditional fractions of low, middle, and high income earners.

Figure 4: Labor Income Calibration

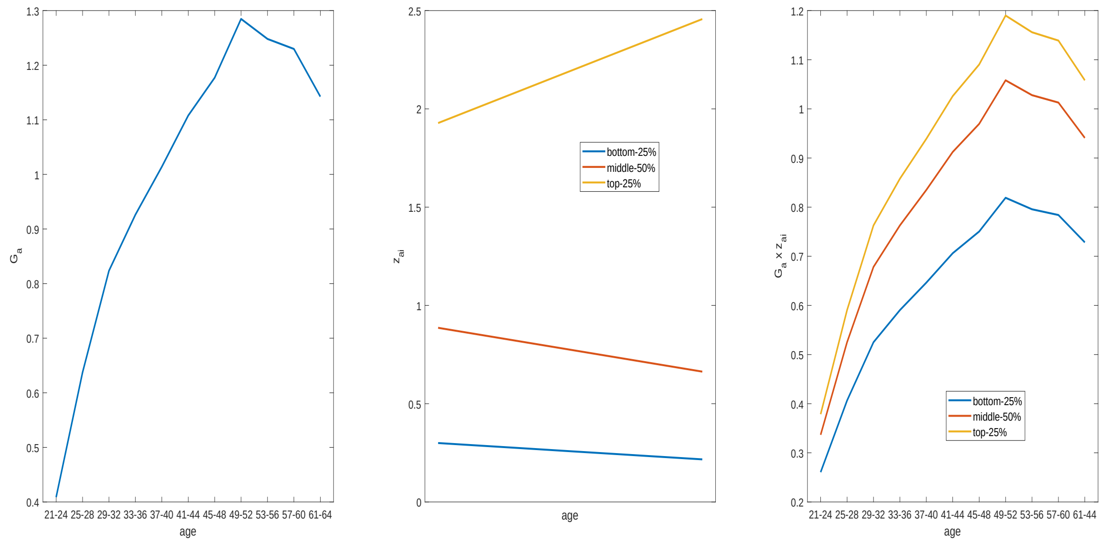

The left panel of Figure 4 plots the age-earnings profile G^a. It shows the familiar hump shape over the life-cycle; G^a is set to zero for age 65 and older. The middle panel plots the idiosyncratic productivity grid at each age: $z_{a,low}$, $z_{a,middle}$, and $z_{a,high}$. It shows that at young ages, dispersion between the earnings groups muted, since the values are all close to 1. The lines diverge with age, implying a rising variance of earnings over the life-cycle. The right panel plots the product, $G^a z_{a,j}$, which are the efficiency units of labor for the various productivity groups.

In the model, we are calibrating each household's productivity. In the data, we do not observe productivity, but use earnings as our productivity estimate. If hours worked differed drastically across age and income groups, then the model's income distribution might look very different from that of the data. Hours worked are relatively flat across productivity and age, both in the model and in the data. Therefore, the model's (endogenous) income distribution looks similar to the data, as can be seen from the top panels of Figure 2.

C 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 as:

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

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 (4) delivers:

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

Rewrite in logs, using lowercase letters to denote logs:

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

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

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

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

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

Invert this expression delivers the housing supply elasticity:

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

If the elasticity of wages to housing supply is small ($\frac{\partial w}{\partial h} \approx 0$), 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. If H is far below $\overline{H}$, that term is close to zero and the housing supply elasticity is close to $\frac{\rho_h}{1-\rho_h}$.

In the baseline calibration, $\rho_h = .66$, $H/\overline{H} = 0.133$ and $\frac{\partial w}{\partial h} = 0.28$. The latter is calculated as the percentage change in wages divided by the percentage change in housing. These numbers deliver a house price elasticity of 0.97, close to the target of 1.0.

D Data Appendix

D.1 The Typical U.S. Metropolitan Area

We compile data on the largest 75 metropolitan statistical areas for the year 2016. From the Census Bureau, we collect population for each MSA and each county. We also obtain data on the number of residents that are above age 65 and the number of residents that are above 21. Their ratio is the share of retirees. Finally, we obtain counts of the number of owner- and renter-occupied housing units, allowing us to construct the home ownership rate.

From Zillow, we collect the Zillow House Value Index (ZHVI) and the Zillow Rental Index (ZRI) for each MSA and each county. The ZHVI is the price of a typical dwelling (single-family, condo, coop), while the ZRI is the monthly rent for a typical house (single-family, apartment). Zillow uses a machine-learning algorithm to control for property characteristics so that the resulting

ZHVI and ZRI indices pertain to the *same property* in a given period and location.

We match Census and Zillow data at the county level, which requires some manual work because of different naming conventions. The 75 MSAs are comprised of 487 counties.

We then compute for all 75 MSAs the following 6 statistics: (1) MSA population, (2) MSA ZHVI P, (3) MSA ZRI R, (4) MSA price to annual rent ratio P/R, (5) the ratio of the msa population that is above age 65 to the msa population above age 21 ($fracret$), (6) the home ownership rate in the MSA (Own). Finally, we calculate the population-weighted and equally-weighted averages of these moments among the 75 MSAs to obtain statistics for the "typical MSA." Table 5 shows these statistics for the largest 40 MSAs by population as well as the population-weighted average among all 75 MSAs.

D.2 Out-of-town Housing Demand

Out-of-town (OOT) housing demand is estimated using the data set compiled for us by `CoreLogic`. This confidential data set contains the monthly time series of number of housing purchases for Manhattan and for NYC MSA between January 2004 and September 2016. Housing purchases are defined as purchases of single-family, 2-4 family, condominiums, and co-ops. OOT purchases are identified using the reported mailing addresses on payment/tax forms. Specifically, if the address of a buyer is either abroad or not contained in the list of 1,304 ZIP codes inside NYC MSA, then the transaction is classified as an OOT purchase.

One complication arises because not only individuals but also companies purchase residential real estate. We include purchases by the following types of corporate entities: LLC, Inc, Corp, and Trust. Combined, these account for 7.28% of all purchases in the New York metro and even 11.13% in Manhattan. We have an address for these corporate purchases as well. Following the same address rules, we obtain the number of OOT corporate purchases and the number of NY MSA corporate purchases in each month. If the buyer of an apartment is a corporation, we can not be certain whether the individual who ultimately owns the apartment is a local or from OOT. Some OOT corporate purchases may be done by locals while some NY MSA corporate purchases may actually hide the identity of OOT investors. Under *assumption 1*, we assume that all OOT corporate purchases are by OOT investors and none of the NY MSA corporate purchases are by OOT buyers. Under *assumption 2*, we assume that 70% of all OOT corporate purchases are by OOT investors and 30% of the NY MSA corporate purchases are by OOT buyers. We have also computed the OOT share assuming (90%,10%) and (80%,20%) assumptions and the results are in between those for assumption 1 and assumption 2. Since there are a lot more NY MSA corporate purchases than OOT corporate purchases, the OOT share under assumption 2 is higher than under assumption 1.

As described in Table 6, the average OOT purchase fractions are 9.2%/2.8% for Manhattan/Rest of NY metro under *assumption 1*, while the fractions are 11.6%/4.6% under *assumption 2*. Based on conversations with market participants, we believe assumption 2 comes closer to approximating the true OOT share. A full-sample OOT share of 5% for the entire NYC metro area and of 10% for Manhattan seem conservative (Panel C).

D.3 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).

Table 5: U.S. Metropolitan Area Statistics

msa name	Pop	P	R	P/R	fracret	Own
New York	20,153,634	498,955	2,404	17.30	19.1%	51.5%
Los Angeles-Long Beach-Anaheim	13,310,447	557,640	2,587	17.96	17.1%	48.4%
Chicago	9,512,999	200,328	1,655	10.09	17.7%	64.2%
Dallas-Fort Worth	7,233,323	184,461	1,534	10.02	14.8%	59.7%
Houston	6,690,766	179,457	1,575	9.49	14.2%	60.1%
Washington	6,131,977	398,394	2,194	15.13	15.6%	63.0%
Philadelphia	6,070,500	203,562	1,590	10.67	19.7%	67.4%
Miami-Fort Lauderdale	6,066,387	231,556	1,889	10.21	22.5%	60.1%
Atlanta	5,789,700	178,136	1,344	11.04	15.2%	63.0%
Boston	4,794,447	410,340	2,296	14.90	19.2%	61.3%
San Francisco	4,679,166	829,218	3,389	20.39	18.3%	53.7%
Phoenix	4,661,537	222,821	1,313	14.15	20.0%	61.4%
Riverside	4,527,837	308,608	1,704	15.09	17.3%	61.9%
Detroit	4,297,617	128,608	1,192	8.99	20.0%	68.6%
Seattle	3,798,902	418,325	2,072	16.83	16.5%	59.8%
Minneapolis-St Paul	3,551,036	229,230	1,558	12.26	17.1%	69.7%
San Diego	3,317,749	509,700	2,431	17.47	17.4%	52.7%
Tampa	3,032,171	165,687	1,329	10.39	24.6%	63.9%
Denver	2,853,077	341,064	2,016	14.10	15.9%	63.3%
Baltimore	2,798,886	240,402	1,766	11.35	19.0%	66.0%
St. Louis	2,785,284	143,490	1,140	10.49	20.1%	69.0%
Charlotte	2,442,133	167,584	1,249	11.18	17.2%	65.3%
Orlando	2,441,257	188,415	1,380	11.38	18.8%	60.2%
Portland	2,424,955	342,224	1,774	16.08	17.9%	60.8%
San Antonio	2,380,812	163,058	1,320	10.30	17.4%	61.8%
Sacramento	2,296,418	343,465	1,688	16.96	19.3%	59.0%
Las Vegas	2,155,664	208,300	1,241	13.99	18.3%	52.3%
Pittsburgh	2,155,452	122,317	1,107	9.21	24.0%	69.2%
Cincinnati	2,142,179	140,438	1,263	9.27	19.0%	65.8%
Austin	2,056,405	266,196	1,735	12.79	13.3%	58.1%
Cleveland	2,055,612	122,949	1,163	8.81	22.4%	65.0%
Columbus	2,041,520	144,022	1,330	9.03	16.7%	61.1%
Kansas City	2,028,527	158,908	1,264	10.48	18.4%	64.7%
Indianapolis	2,004,230	139,160	1,211	9.58	17.4%	64.7%
San Jose	1,978,816	945,350	3,496	22.53	16.7%	56.7%
Nashville	1,865,298	211,455	1,542	11.43	16.7%	65.3%
Virginia Beach	1,726,907	212,569	1,433	12.36	17.9%	61.0%
Providence	1,614,750	243,116	1,604	12.63	21.1%	60.5%
Milwaukee	1,572,482	178,678	1,302	11.43	19.2%	60.0%
Jacksonville	1,478,212	168,442	1,304	10.77	19.3%	64.5%
Pop-weighted avg 75 MSAs	6,075,613	297,875	1,775	13.03	18.3%	60.6%
Equally-weighted avg 75 MSAs	2,650,526	238,508	1,551	11.97	19.0%	63.1%

Table 6: Fraction of OOT Purchases of New York Housing Units

Panel A: Manhattan				
Assumption	**2004.01-2016.09**	2004.01 - 2007.12	2008.01 - 2011.12	2012.01 - 2016.09
1	9.2%	8.2%	9.0%	10.1%
2	11.6%	9.6%	11.4%	13.4%
Panel B: Rest of New York metro area				
1	2.8%	2.5%	2.6%	3.1%
2	4.6%	3.6%	4.4%	5.8%
Panel C: New York metro area				
1	3.3%	2.9%	3.3%	3.8%
2	5.3%	4.0%	5.0%	6.5%

Notes: Share of residential real estate purchases made by out-of-town (OOT) buyers in the New York metropolitan area. Source: Core Logic. Monthly data from January 2004 through September 2016.

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 other than Manhattan and call them 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 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.