Mortgage Interest Deductions? Not a Bad Idea After All

Shahar Rotberg[*] Joseph B. Steinberg[†]

First Draft: February 2023
This Draft: February 2023

Abstract

Previous studies have found that homeownership subsidies like the mortgage interest deduction (MID) reduce aggregate welfare. We argue that these studies have failed to account for two important features of the rental market: the price elasticity of the rental supply, which determines how much rents change when rental demand shifts; and the number of renters who spend most of their income on housing, which determines how much changes in rents affect welfare. We develop a method to estimate the rental supply elasticity and show that when both features are accounted for, the MID and other homeownership subsidies actually increase welfare.

JEL Classification: R21, R31, H20.

Keywords: Mortgage interest deduction, rental market, imputed rents, homeownership, welfare

1 Introduction

Promoting homeownership is a cornerstone of U.S. economic policy. The government subsidizes homeowners in a variety of ways, from guaranteeing mortgage-backed securities to providing a tax deduction for mortgage interest payments. In equilibrium, these policies also affect renters because lowering the cost of homeownership reduces demand for rental housing. We argue that previous studies have underestimated the importance of these effects by failing to properly account for two key features of the rental market: (i) the price elasticity of the supply of rental housing, which determines how much rents

[*]shahar.rotberg@mail.utoronto.ca, University of Toronto

[†]joseph.steinberg@utoronto.ca, University of Toronto. I acknowledge the Social Sciences and Humanities Research Council of Canada (grant number 506633) for financial support.

change when demand shifts; and (ii) the number of renters who spend the majority of their income on housing, which determines how many households are severely affected by changes in rents. We show that once these features are properly accounted for, home-owner subsidies that previous studies have found to be harmful are actually beneficial, especially for low-income renters.

The consensus in the quantitative housing literature is that subsidizing homeowners reduces aggregate welfare. The basic logic behind this consensus is that homeowners are richer than renters on average, so resources should be redistributed from the former to the latter, not the other way around. Two policies that have been studied extensively are the mortgage interest deduction (MID) and the preferential tax treatment of imputed rents generated by owner-occupied housing. Gervais (2002), Chambers et al. (2009), Floettoto et al. (2016), and Nakajima (2020) all find that eliminating both of these subsidies would increase welfare, and further work by Sommer and Sullivan (2017) and Karlman et al. (2021) report welfare gains from eliminating the MID.[1] However, we argue that none of these studies accurately model the effects of these reforms on the rental market. In some, the price elasticity of the supply of rental housing is higher than in the data, which means that rents respond too little in equilibrium to changes in rental demand. In others, fewer renters spend the majority of their income on housing than in the data, which means that changes in rents have too little effect on renters' welfare. Table 1 provides an overview of these studies, their assumptions, and their findings.

We calibrate a dynamic general equilibrium model to match empirical evidence on the rental supply elasticity and the fraction of renters who spend at least half of their income on housing and find that eliminating the MID and taxing imputed rents would actually reduce welfare, not increase it. Both pieces of evidence are crucial: when our model is calibrated to match either one but not the other, these reforms would increase welfare as previous studies have found.

The elasticity of the rental housing supply is important because it determines how much rents change in equilibrium when demand for rental housing shifts, but to our

[1] Floettoto et al. (2016) and Karlman et al. (2021) report short-run welfare losses in some experiments. Our focus is on the long run, where all of these studies agree.

knowledge there are no existing estimates of this elasticity in the literature. We propose a new approach to identify this key parameter in structural models using evidence on how rents respond to changes in property taxes. The core idea behind this approach is that the more elastic the supply, the greater the share of property taxes borne by renters. The Institute on Taxation and Economic Policy (ITEP) estimates that renters pay about 50% of property taxes in the United States (Wiehe et al., 2018), which implies a rental supply elasticity of 1.8 in our model. In comparison, many studies in the quantitative housing literature simply assume an infinite supply elasticity (see, e.g., Gervais, 2002; Kaplan et al., 2019; Karlman et al., 2021). Other studies feature models in which rental housing is supplied endogenously by households that choose to act as landlords, but the supply elasticity in these studies still tends to be much higher than in the data (see, e.g., Floettoto et al., 2016; Sommer and Sullivan, 2017).[2] The high rental supply elasticities typically seen in the literature imply that removing the MID or taxing imputed rents would increase rents far less than the evidence on property tax incidence indicates.

The number of renters who spend the majority of their income on rent—the definition of "severely cost-burdened" according to the U.S. Department of Housing and Urban Development (HUD)—is important because it determines the welfare consequences of changes in rents. Using data from the 2017 American Housing Survey (AHS), we find that 15.1% of renters in the United States are in this group. We use this moment to identify the minimum allowable size of a rental dwelling in our model, which we find to be about 29% of the average owner-occupied house size. In comparison, many quantitative studies of housing policy have no lower bound on rental housing size (see, e.g., Chambers et al., 2009; Floettoto et al., 2016; Nakajima, 2020). In these models, renters spend less of their income on housing and can respond to rent increases by downsizing further, which implies that changes in rent have smaller welfare effects. Other studies do not report how this parameter is chosen at all (see, e.g. Sommer and Sullivan, 2017; Karlman et al., 2021). Of course, the minimum rental size is a real-world constraint imposed by many

[2]The sole exception is Chambers et al. (2009), whose model features a realistically low supply elasticity but still predicts that removing the MID would be optimal because it does not match the number of cost-burdened renters. Our study shows why matching both the supply elasticity and the share of cost-burdened renters is crucial for accurately quantifying the consequences of housing policy.

jurisdictions, not just a model ingredient that is useful for generating a realistic rent-to-income distribution. For example, New York City zoning law requires rental apartments to be at least 300 square feet.[3]

Our model is populated by overlapping generations of households that are heterogeneous in labor productivity, financial wealth, housing tenure, and mortgage debt. Households supply labor inelastically and are taxed progressively on their earnings as in Heathcote et al. (2017). They obtain housing services by renting or buying a home. Rental housing is provided by a competitive rental company as in Chambers et al. (2009). Home purchases can be financed with long-term mortgages. Households can deduct mortgage interest payments from their taxable income as allowed by the U.S. tax code. Most importantly for our purposes, there is a minimum allowable rental unit size and the cost of managing rentals is convex. We calibrate our model so that it reproduces salient facts about demographics, income dispersion, public finances, and housing markets. Several of our housing-related target statistics are standard in the quantitative housing literature: the maximum leverage ratio; the mortgage interest premium; property taxes and transaction costs; the homeownership rate; and the rent-price ratio. We add two new targets that capture the key features of the rental market discussed above: the share of property taxes borne by renters, which identifies the convexity of the rental management cost and determines the supply price elasticity; and the share of renters who are cost-burdened, which identifies the minimum rental size.

We use our model to analyze two widely-studied policy changes: repealing the MID and taxing the imputed rents generated by owner-occupied housing. In order to assess the role of the supply elasticity and the share of cost-burdened renters, we conduct each experiment in three different calibrations of our model: the baseline calibration that matches the evidence on both features of the rental market structure; a calibration with an infinitely-elastic rental supply that still matches the number of cost-burdened renters; and a calibration without a lower bound on rental size that still matches the property tax incidence of renters. In our baseline calibration, both policy changes would increase rents and reduce welfare. In the high-elasticity calibration, both policies would increase

[3]See https://zr.planning.nyc.gov/article-ii/chapter-3/23-23.

welfare because they would reduce rents. In this case, equilibrium rent is independent of the quantity of rental housing supplied, so the decrease in home prices caused by removing homeownership subsidies would cause rents to fall. In the no-minimum-rental calibration, both policies would increase welfare, but for different reasons. Removing the MID would actually increase homeownership slightly, consistent with previous studies that use a similar calibration such as Chambers et al. (2009), Sommer and Sullivan (2017), Karlman et al. (2021). This would cause rents to fall, benefiting renters. On the other hand, taxing imputed rents would cause rents to rise in this calibration, but this would have a smaller negative welfare effect than in the baseline calibration because fewer renters are cost-burdened, and thus the benefits of the policy would outweigh these losses.

We make two methodological contributions in addition to our quantitative contributions about the impact of homeownership subsidies. First, we demonstrate how to identify the elasticity of the rental supply in structural models using data on the share of property taxes borne by renters. No previous studies have used data to discipline this elasticity, and our findings show that this has contributed to mis-measurement of the effects of housing policy. Second, we propose a novel approach to modeling long-term mortgage debt that features origination costs, borrowing constraints that apply only at origination, and repayment constraints that apply throughout the life of a mortgage. It generates payment schedules that closely match up with the actual payment schedule for a 30-year amortization period, but is still just as tractable as previous approaches like Karlman et al. (2021) and Chatterjee and Eyigungor (2015) that treat mortgages as perpetuities with geometrically-decaying coupons. Both of these innovations have the potential to improve quantitative analysis of many housing-related issues in the future.

2 Key factors: empirical evidence and identification

The main finding of this paper is that there are two key factors that determine the welfare consequences of homeownership policies: the price elasticity of the rental supply curve and the number of renters who spend most of their income on housing. Here, we discuss the empirical evidence on these factors, how we use this evidence to identify structural parameters in our model, and how our calibration compares to other studies.

2.1 Rental supply elasticity

The first factor that determines the consequences of homeownership policies is the price elasticity of the supply of rental housing. The steeper the slope of the rental supply curve, the more rents change in equilibrium when demand for rental housing shifts as a result of a change in the homeownership rate. Panel (a) of Figure 1 provides a graphical illustration of this point. The solid red line, S_1, depicts an elastic rental supply curve, and the solid blue line, S_2, an inelastic supply curve. The initial equilibrium price, P_0, is the same under both supply curves, but when demand shifts outward from D to D', the equilibrium price rises more under the inelastic supply curve than the elastic one.

We propose a method to identify the rental supply elasticity using empirical evidence on the share of property taxes borne by renters. The gist of our identification strategy is that the slope of the rental supply curve determines how much rents change in equilibrium when the property tax rate changes.[4] If the supply were perfectly inelastic, rents would be independent of property taxes, and so renters' share of property taxes would be zero. At the other extreme, if the rental supply were perfectly elastic, rents would change one-for-one with the property tax rate, i.e., renters' share of property taxes would be 100%. In between these two extremes, the slope of the supply curve determines the extent to which changes in property taxes are passed on to renters.

Panel (b) of Figure 1 provides a graphical illustration of this identification. As before, the solid red and blue lines, S_1 and S_2, show elastic and inelastic rental supply curves, respectively, and the initial equilibrium price, P_0, is the same in both cases. The dashed red and blue lines, S_1' and S_2', depict the elastic and inelastic supply curves after a property tax increase. The equilibrium price after the property tax increase in the elastic case is labeled P_1', and P_2' is the equilibrium price after the tax increase in the inelastic case. The share of the property tax increase borne by renters is calculated as the change in the equilibrium price divided by the vertical distance between the supply curve before and after the shift. The denominator in this expression represents the difference between the rental supplier's pre- and post-tax revenue. Because both supply curves shift upward

[4] The first-order condition (17) of the rental management company in our model provides a mathematical formalization of this logic.

by the same amount—the size of the tax increase—it is clear from the fact that $P'_1 > P'_2$ that renters' tax incidence is higher under the elastic supply curve. To calibrate the rental supply elasticity in our model, we simulate a 1 percentage point increase in taxes on the value of rental property and compute the share of that increase borne by renters, and set the convexity of the rental management cost function so that this share is equal to its empirical counterpart.

Panel (a) of Table 2 provides an overview of the empirical evidence on renters' share of property taxes. The Institute on Taxation and Economic Policy (ITEP), whose analysis of tax burden distribution in the United States is widely relied upon by policy analysts, estimates a value of 50% (Wiehe et al., 2018). Recent studies that use property-level panel data like Fack (2006) and Schwegman and Yinger (2020) report lower numbers ranging from 15–25%. Earlier studies that use region-level data like Orr (1970) typically report shares of 40–50% closer to the ITEP estimate, although England (2016) and Schwegman and Yinger (2020) argue that these estimates suffer from a variety of mis-specification issues such as simultaneity bias. We use the ITEP estimate of 50% because it lies toward the upper end of the range of figures reported in the literature, which implies a conservatively high rental supply elasticity of 1.8 in our model. Our findings imply that policies like removing the MID would have significantly worse welfare consequences if the rental supply was less elastic as implied by some of the more recent contributions to this literature.

Additional supporting evidence comes from studies on the extent to which rents respond to the provision of housing vouchers. The basic idea here is that, all else being equal, the more elastic the rental supply curve, the less rents should respond in equilibrium to a shift in rental demand caused by vouchers. Panel (b) of Table 2 provides an overview of empirical estimates of the effect of vouchers on rents. Although some early studies report small effects (see, e.g., Barnett, 1979), most recent studies like Susin (2002), Gibbons and Manning (2006), Kangasharju (2010), Viren (2013), Collinson and Ganong (2018), and Sayag and Zusman (2020) estimate that landlords capture significant portions of voucher values. We view this as supporting evidence that the rental supply is far from perfectly elastic. It does not, however, provide direct evidence on the magnitude of the rental supply elasticity. Given an estimate of the effect of vouchers on rent, one would need additional information such as the income elasticity of demand and the voucher value relative to the average renter's income to infer the supply price elasticity.

Previous quantitative studies have taken two approaches to modeling the supply of rental housing. The first is to model a representative rental management corporation. This approach has been used in a number of studies of the MID (see, e.g., Gervais, 2002; Nakajima, 2020; Karlman et al., 2021), as well as studies of other housing-related phenomena (see, e.g., Arslan et al., 2015; Kaplan et al., 2019). Almost universally, these studies assume that the rental supply is infinitely elastic. In the second approach, rentals are supplied by individual households that choose to act as landlords (see, e.g., Sommer et al., 2013; Floettoto et al., 2016; Sommer and Sullivan, 2017). The rental supply elasticity in these endogenous-landlord models is endogenous but tends to be substantially higher than our estimate. For example, the results reported by Floettoto et al. (2016) imply an elasticity of 38, which corresponds to a property tax incidence for renters of more than 90% .[5] In short, the rental supply is substantially more elastic under both approaches than indicated by the evidence on property tax incidence, which implies that shifts in demand for rental housing lead to counterfactually small changes in rents.

The one exception is Chambers et al. (2009), which is a mixture of these two approaches. It features a representative rental corporation with a finitely-elastic supply curve and individual households that supply rentals endogenously, resulting in an aggregate supply elasticity of about 1.6. However, this study does not identify the corporate supply elasticity in their model using any empirical evidence (they simply set it arbitrarily); providing a means of identifying this crucial parameter is one of our methodological contributions. Moreover, despite the fact that this study has a realistically low aggregate rental supply elasticity, it incorrectly predicts that removing the MID would be welfare-improving because it does not account for our second key factor: many renters spend large shares of their income on rent.[6]

[5]We study a version of the endogenous-landlord model in the appendix. We confirm that the aggregate rental supply in this model is highly elastic and that the effect of homeownership policies in this model are similar to the effects in a model with an infinitely-elastic rental management company.

[6]Chambers et al. (2009) impose no restriction on the minimum size of a rental dwelling and this implies that all renters spend the same amount on housing relative to goods consumption. When the MID is

2.2 Prevalence of cost-burdened renters

The second factor is that many renters spend large shares of their income on housing. For these cost-burdened households, even small changes in rents have large welfare consequences, and the more of these households there are, the more important these consequences are for aggregate welfare.

Figure 2 shows the distribution of housing costs relative to income for renters in the 2017 American Housing Survey (AHS). We include all forms of income in the AHS data, including government transfers which are important for many households with low market incomes. We use the HUD's definition of a severely cost-burdened household as one that spends at least 50% of its gross income on housing (Larrimore and Schuetz, 2017).[7] In the AHS data, 15.1% of U.S. renters are severely cost-burdened. Of course, some renters spend substantially larger shares of their income on housing; for example, almost 5% of renters spend more than 75%.

In our model, the number of cost-burdened renters is governed by the minimum size of a rental dwelling. The identification here is straightforward: low-income households are constrained by the minimum rental size, and when this size increases they are forced to spend a larger fraction of their income on rent. To match the share of cost-burdened renters in the data, our model requires a minimum rental size that is about 29% of the average owner-occupied house size. In section 4.3, we show that our model matches the entire distribution of rent-to-income closely, especially on the right tail which is most important for capturing the welfare consequences of changes in rents.

There is little consistency in how previous quantitative studies have disciplined the

removed in their model, rents rise by 4%, which would have massively negative welfare consequences for low-income renters if they matched the share of cost-burdened renters as we do. In our model, rents rise by only 0.6% and this is enough to generate large welfare losses.

[7] The HUD defines a household that spends at least 30% of its income on housing as cost-burdened, and a household that spends at least 50% as severely cost-burdened. The former definition is more widely used to assess housing affordability in the United States (Herbert et al., 2018). We use the latter definition as our calibration target because it allows us to do a significantly better job of matching the entire right tail of the rent-to-income distribution. Calibrating to the 30% threshold instead leads our model to substantially understate the number of households who spend larger shares of their income on rent, whereas targeting the 50% threshold allows us to also match closely the number of renters who spend 60%, 75%, 80%, etc. of their income on rent, which are important for welfare outcomes. For brevity, we use the terms "cost-burdened" and "severely cost-burdened" interchangeably.

minimum rental size. Some papers, like Sommer et al. (2013), Sommer and Sullivan (2017), and Karlman et al. (2021), do not report values or calibration targets for this parameter. Others like Chambers et al. (2009) and Floettoto et al. (2016) assume there is no lower bound on the size of rental dwellings at all. This explains why Chambers et al. (2009) find that repealing the MID would improve welfare even though their rental supply elasticity is close to ours. The closest calibration to ours on this dimension is Gervais (2002), who chooses the minimum rental size so that households in the bottom quintile of the income distribution devote 50% of their spending to housing on average. This study still finds that repealing the MID would be optimal because it assumes an infinitely elastic rental supply. As we show in our quantitative analysis below, repealing the MID lowers welfare only when one takes both of the key factors discussed in this section into account.

3 Model

The model economy features overlapping generations of finitely-lived households, a construction company, a rental company, and a government. Households supply labor inelastically and choose consumption of goods and housing services; whether to rent, buy, or sell a home; how much to save; and how much mortgage debt to take on. The construction company builds new housing and the rental company purchases existing housing units to rent out. The government levies taxes on income and property and provides social security benefits to retirees. The key ingredients that help us match the two key facts described above are the convexity of the rental company's management cost, which governs the slope of the rental supply curve; and the minimum size of a rental dwelling, which determines the number of renters that spend the majority of their income on housing.

In many respects, our model is similar to others in the quantitative housing literature, particularly those listed in Table 1. Like Karlman et al. (2021), we include a detailed demographic structure that realistically captures the full household life cycle. Like Chambers et al. (2009), Sommer and Sullivan (2017), and Karlman et al. (2021), we include a progressive income tax system, which is important for capturing the fact that the benefits of the MID accrue primarily to high-income households. Like Sommer and Sullivan (2017), we allow for a (mildly) price-inelastic aggregate housing stock. Like Gervais (2002), Nakajima (2020), and Karlman et al. (2021), rental housing is supplied only by a competitive rental company, but we allow the rental management cost to be convex as in Chambers et al. (2009).[8] Like Gervais (2002), Sommer and Sullivan (2017), and Karlman et al. (2021), there is a tenure-specific lower bound on house size.

One key difference between our model and those of previous studies is that we introduce a novel method of modeling long-term mortgages that incorporates origination costs, prepayment penalties, and two forms of borrowing restrictions that apply only at origination. This method is no less tractable than previous approaches such as Chatterjee and Eyigungor (2015) and Karlman et al. (2021) that treat mortgages as perpetuities with geometrically decaying coupons, but is flexible enough to generate realistic payment schedules.

3.1 Households

3.1.1 Demographics and preferences

Age is indexed by j. Households are born at age $j = 1$, retire at age $j = J_R$, and live until age $j = J$, after which they die with certainty. Households have preferences over consumption of goods, c, and housing services, h, given by the flow utility function

$$u_j(c, h) = \xi_j^{\sigma} \frac{\left(c^{1-\gamma} h^{\gamma}\right)^{1-\sigma}}{1-\sigma}. \tag{1}$$

The parameter γ is the weight on housing services, σ governs risk aversion and the intertemporal elasticity of consumption, and ξ_j is an equivalence-scale parameter that captures variation in household size over the life cycle.

3.1.2 Endowments

Working-age households ($j < J_R$) receive an idiosyncratic endowment of labor income,

$$y_j(z) = \zeta_j z, \tag{2}$$

[8] We show in the appendix that the approach taken by Floettoto et al. (2016) and Sommer and Sullivan (2017) in which rental housing is supplied by individual households who choose endogenously to be landlords is similar to a model with an infinitely price-elastic rental company.

where ζ_j is a common life-cycle profile and $z \in \mathcal{Z}$ is an idiosyncratic shock that follows a Markov process $F(z, z')$. Retired households ($j \geq J_R$) receive social security benefits given by

$$y_R(z) = \kappa(z)\mathbb{E}[y_j(z)]. \tag{3}$$

The function $\kappa(z)$ represents the fraction of the average labor income that a retiree with idiosyncratic shock z receives. This allows us to capture, in a reduced-form way, the correlation between lifetime labor income and social security benefits.

3.1.3 Housing

Housing services are obtained by renting or owning a home. We denote a household's tenure by $o \in \{0, 1\}$, where 0 indicates a renter and 1 indicates a homeowner. Each unit of housing provides the same amount of housing services regardless of tenure, but there is a tenure-specific lower bound on house size: rental units cannot be smaller than $\underline{h}_r$, while owner-occupied dwellings cannot be smaller than $\underline{h}_o$. The per-unit price of housing is p and the rental price is p_r. Buying and selling a house require one-time proportional transaction costs, τ_b, and τ_s, respectively. Rental and owner-occupied housing are both taxed at a rate τ_p and depreciate in value at a rate δ.

3.1.4 Saving and mortgages

Households can save at an exogenous interest rate of r.[9] We denote a household's savings as a. Unsecured borrowing is not allowed, but homeowners have access to mortgage debt, m. The interest rate on mortgages is $r_m = r + \Delta_m$, where Δ_m is the mortgage premium. Mortgages are modeled as as long-term debt with origination costs, borrowing restrictions, and prepayment penalties.

Originating a new mortgage incurs a fixed cost ω_1. Two restrictions apply at the time

[9]The assumption of a small open economy is prevalent in the quantitative housing literature. See, for example, Sommer et al. (2013), Arslan et al. (2015), Floettoto et al. (2016), Sommer and Sullivan (2017), Kaplan et al. (2019), and Karlman et al. (2021). Under the hood, so to speak, we allow households to save in bonds and also purchase stock in the rental and construction companies, but a no-arbitrage condition requires households to be indifferent between these two options, i.e., the rate of return on stock must be the same as the exogenous world interest rate. Because we do not require the asset market to clear, we do not need to explicitly keep track of the dividends associated with these stock holdings. An alternative interpretation is that the rental and construction companies are owned by foreigners as in Floettoto et al. (2016).

of origination. The first is a loan-to-value (LTV) constraint, which limits mortgages to a fraction λ_1 of the house's purchase price ph':

$$m' \leq \lambda_1 ph'. \tag{4}$$

The second is a gross debt servicing (GDS) constraint, which limits carrying costs to a fraction λ_2 of gross income:

$$\tau_p ph' + (r_m + \nu_1)m' \leq \lambda_2[y_j(z) + ra], \tag{5}$$

where ν_1 is the average fraction of principal that mortgage holders pay annually over the life of a mortgage. Retirees cannot originate new mortgages but can carry existing debt into retirement.

Homeowners with existing mortgages originated in previous periods are required to pay at least a fraction ν_2 of their mortgage debt, and are allowed to pay up to a fraction $\nu_3 > \nu_2$ without penalty. Any payment above this amount triggers a repayment penalty given by

$$\tau_{pp}(m, m') = \omega_2 \max\{(1 - \nu_3)m - m', 0\}. \tag{6}$$

When mortgage debt falls below a fraction ν_4 of the house value (i.e. $m \leq \nu_4 ph$), the mortgage holder is required to either repay the remainder of the debt without penalty or originate a new mortgage.

Our approach to modeling long-term mortgages differs from other approaches in the literature. Many studies such as Chatterjee and Eyigungor (2015) and Karlman et al. (2021) model mortgages as perpetuities with geometrically decaying coupons. Their approach is tractable because it does not require one to keep track of additional state variables aside from the mortgage balance, but it implies that mortgages are repaid significantly more slowly than in the data. Others like Chambers et al. (2009) allow for mortgages that are amortized over a finite number of periods, but this comes at a significant computational cost because it requires one to keep track of the number of payments remaining on the mortgage as well as the original balance. Our approach, on the other hand, is flexible enough to generate realistic repayment trajectories by appropriately calibrating the parameters ν_1, ν_2, ν_3, and ν_4, but is also highly tractable because the current mortgage balance encodes all the relevant information about the household's future repayment schedule.

3.1.5 Income taxes

Working-age households pay taxes on their labor income after taking a personal exemption and either a standard or itemized deduction. Taxable labor income is given by

$$\tilde{y}_j(z, o, h, m) = \max\{y_j(z) - \tau_e - \max[\tau_d, o(\tau_m r_m m + \tau_p p h)], 0\}, \tag{7}$$

where τ_e is the personal exemption, τ_d is the standard deduction, and τ_m is the fraction of mortgage interest that can be deducted if a household chooses to itemize its mortgage interest payments. Labor income taxes are progressive following Heathcote et al. (2017). A household's labor income tax liability is

$$\tau_j(z, o, h, m) = \tilde{y}_j(z, o, h, m) - \tau_\ell \tilde{y}_j(z, o, h, m)^\psi, \tag{8}$$

where τ_ℓ controls the average level of taxation in the economy and ψ controls the degree of progressivity. All households pay a proportional tax τ_k on capital income.

3.1.6 Dynamic program

A household's state variables are its idiosyncratic labor income, z, its financial wealth, a, its tenure, o, its house size, h, and its mortgage debt, m. Given these state variables, which we collectively denote by s, it chooses how much to save, a'; whether to change its tenure, o'; how much goods and housing services to consume, c and h'; whether to originate a new mortgage, n; and how much of its existing mortgage debt to repay or how much new mortgage debt to take on, m'. The dynamic program that represents this optimization problem is

$$V_j(s) = \max_{c, a', o', h', n, m'} \left\{ u_j(c, h') + \beta \int_{\mathcal{Z}} V_{j+1}(s') \, \mathrm{d}F(z, z') \right\} \tag{9}$$

subject to

$$c + a' + r_m m + (1-o')p^r h + o\left[\delta + \tau_p + \mathbb{1}_{\{o'=0 \vee h' \neq h\}}\tau_s\right] ph + o'\mathbb{1}_{\{o=0 \vee h' \neq h\}}(1+\tau_b)ph'$$
$$= y_j(z) - \tau_j(z, o, h, m) + [1 + r(1-\tau_k)]a + m' - m - n\omega_1 - (1-n)\tau_{pp}(m, m') \tag{10}$$
$$a' \geq 0, \tag{11}$$
$$h' \geq o'\underline{h}_o + (1-o')\underline{h}_r \tag{12}$$
$$n \in \{0, o'\mathbb{1}_{\{j<j_R\}}\} \tag{13}$$
$$nm' \leq \lambda_1 ph' \tag{14}$$
$$n[\tau_p ph' + (r_m + \nu_1)m'] \leq \lambda_2(y_j(z) + ra) \tag{15}$$
$$(1-n)m' \leq \mathbb{1}_{\{m \geq \nu_4 ph\}}(1-\nu_2)m \tag{16}$$

The first constraint (10) is the budget constraint. The left-hand side includes spending on consumption, saving, mortgage interest, rent (if $o' = 0$), property taxes and depreciation (if $o = 1$), and the costs of buying and/or selling houses, which apply only when tenure and/or house size change. The right-hand side includes after-tax income, financial wealth, funds raised by taking on a mortgage, origination costs for new mortgages (if $n = 1$), and prepayment costs for early payments (if $n = 0$). The second constraint (11) states that unsecured borrowing is not allowed. The third constraint (12) states that renters cannot live in houses smaller than $\underline{h}_r$ and owners cannot live in houses smaller than $\underline{h}_o$. The fourth constraint (13) states that a household can only originate a new mortgage if it is below retirement age and is a homeowner at the end of the period. The fifth and sixth constraints (14)–(15) are the LTV and GDS constraints, which apply only when the household originates a new mortgage. The last constraint (16) states that households with existing mortgages must repay at least a fraction ν_2 of the principal if they owe more than a fraction ν_4 of the house value, and must repay the entire principal they owe if their debt is less than this amount. We denote the policy functions for consumption, savings, tenure, housing, mortgage origination, and mortgage debt by $c_j(s)$, $a'_j(s)$, $o'_j(s)$, $h'_j(s)$, $n_j(s)$, and $m'_j(s)$, respectively.

3.2 Rental company

We use the corporate rental market structure from Chambers et al. (2009). The rental company begins each period with a stock of rental housing, S, on which it pays property taxes and depreciation costs. It chooses how much additional rental housing to purchase, $S' - S$, earns rental income $p_r S'$, and incurs management costs $C_r(S') = \frac{\theta_1}{\theta_2} S'^{\theta_2}$. The rental company's dynamic program is

$$W_r(S) = \max_{S'} \left\{ p_r S' - C_r(S') - p(S' - S) - p(\delta_h + \tau_p)S + \frac{1}{1+r} W_r(S') \right\}.$$

The solution to this problem is given by

$$p_r = \theta_1 (S')^{\theta_2 - 1} + \left\{ \frac{r + \delta_h + \tau_p}{1+r} \right\} p. \tag{17}$$

The parameter θ_2 governs the elasticity of the rental housing supply to the rental price. When $\theta_2 = 1$, the rental supply is perfectly elastic as in Gervais (2002) and Kaplan et al. (2019), which means that the price-rent ratio is independent of the quantity of rental housing. When $\theta_2 > 1$, the rental supply curve slopes upward, which means that the rental price rises in equilibrium when demand for rental housing shifts outward. In our calibration, we choose θ_2 so that the model matches empirical estimates of how the property tax incidence is split between landlords and renters.

3.3 Construction company

As in Sommer and Sullivan (2017) and Rotberg (2022), a construction company chooses how much new housing to build, X, subject to a convex cost $C_x(X) = \frac{\epsilon_1}{\epsilon_2} X^{\epsilon_2}$. The construction company's problem is

$$\Pi_x = \max_X \left\{ pX - C_x(X) \right\}, \tag{18}$$

which yields the following relationship between housing investment and the price of housing:

$$p = \epsilon_1 X^{\epsilon_2 - 1}. \tag{19}$$

This expression shows that the price elasticity of the supply of new housing construction is governed by ϵ_2. The aggregate stock of housing follows the law of motion,

$$H' = (1 - \delta)H + X. \tag{20}$$

In a stationary equilibrium, the aggregate housing stock is constant, i.e., $H' = H = \frac{X}{\delta}$.

3.4 Aggregation and equilibrium

There are two housing market clearing conditions. The first states that total demand for housing from both renters and owners must equal the total supply:

$$\sum_{j=1}^{J} \int_{\mathcal{S}} h'_j(s) \, \mathrm{d}\Psi_j(s) = H', \tag{21}$$

where $\Psi_j(s)$ is the distribution of age-j households over the state space $\mathcal{S}$. The second market clearing condition states that demand for rental housing must equal the quantity supplied by the rental company:

$$\sum_{j=1}^{J} \int_{\mathcal{S}} (1 - o'_j(s))h'_j(s) \, \mathrm{d}\Psi_j(s) = S'. \tag{22}$$

The government's budget also must balance:

$$G + \sum_{j=J_R}^{J} \int_{\mathcal{S}} y_R(z) \mathrm{d}\Psi_j(s) = \sum_{j=1}^{J} \int_{\mathcal{S}} \left[\tau_j(s) + \tau_k ra\right] \mathrm{d}\Psi_j(s) + \tau_p pH. \tag{23}$$

The left-hand side of the government's budget constraint contains expenditures on public goods and social security benefits. The right-hand side includes revenues from income and property taxes. When we calibrate the model to the current U.S. tax code, we set the consumption level of public goods, G, to ensure that this constraint holds as in Díaz and Luengo-Prado (2008).

Households are born as renters with zero financial or housing wealth, so the distribution of newborn households, $\Psi_1(s)$, is given by

$$\Psi_1(S) = \int_{\mathcal{Z}} \mathbb{1}_{\{(z,0,0,0,0)\in S\}} \mathrm{d}\bar{F}(z), \tag{24}$$

where S is a typical subset of the state space and $\bar{F}$ is the unconditional distribution of labor productivity shocks associated with the process $F(z, z')$. The distributions of older households evolve according to the law of motion

$$\Psi_{j+1}(S) = \int_{\mathcal{S}} \left[\int_{\mathcal{Z}} \mathbb{1}_{\left\{(z', a_j'(s), o_j'(z), h_j'(s), m_j'(s))\in S\right\}} \mathrm{d}F(z, z') \right] \mathrm{d}\Psi_j(s). \tag{25}$$

A stationary equilibrium is a collection of aggregate prices and quantities, value and policy functions, and distributions,

$$\left\langle p, p_r, H', S', \left\{ V_j(s), c_j(s), a_j'(s), o_j'(s), h_j'(s), n_j(s), m_j'(s), \Psi_j(s) \right\}_{j=1}^{J} \right\rangle, \tag{26}$$

that solve the household's dynamic problem (9), satisfy the rental company's first-order condition (17), satisfy the construction company's first order condition (19), clear the markets for housing and rentals (21)–(22), satisfy the government's budget constraint (23), and satisfy the distributional laws of motion (24)–(25).

4 Calibration

We calibrate our model so that its stationary equilibrium matches salient features of the U.S. economy under the current tax code. We first assign standard values to common parameters and apply estimates from other studies for parameters that have clear empirical counterparts. We then jointly calibrate the remaining parameters so that the model matches a set of important facts about the housing market. Table 3 lists all the calibrated parameters alongside their source or target moment.

4.1 Assigned parameters

There are several groups of externally assigned parameters. They are listed in panels (a)–(e) of Table 3.

Demographics and preferences. A model period corresponds to one year. Households start their lives at age 26, retire at age 66, and live up to age 85, implying that lifespan $J = 60$ and retirement age $J_R = 41$. We set the utility's relative risk aversion parameter, σ to 2. The equivalence scales are set using data from the 2017 American Housing Survey (AHS). We first compute the average household size by age of the head of household, and then take the square root as is common in the literature. The discount factor, β, and the share of housing services in utility, γ, are determined in the second stage of our calibration procedure.

Labor income process. We assume that labor productivity shocks follow an AR(1) process in logs,

$$\log z' = \rho_z \log z + \epsilon_z, \quad \epsilon_z \sim N(0, \sigma_z), \tag{27}$$

We set $\rho_z = 0.94$ and $\sigma_z = 0.20$ following Rotberg and Steinberg (2022). We set the deterministic life-cycle component of labor market productivity to $\zeta_j = 1 + \min\{0.38(j - 1)/30, 0.38\}$ so that households' labor income rises by 38 percent by age 55 (Guvenen et al., 2019a). We follow Guvenen et al. (2019b) in setting $\kappa(z)$, which determines the extent to which social security benefits depend on productivity shocks.

Taxes. The interest rate is set to 3% as in Kaplan et al. (2019) and the capital income tax rate is set to 15 percent as in Karlman et al. (2021). We set the personal exemption, τ_e, to 43% of the standard deduction (the equivalent of $4,050 as in the U.S. tax code). We set the labor income tax progressivity parameter, ψ, to 0.85 based on Heathcote et al. (2017). The average level of taxation, τ_ℓ, and the standard deduction, τ_d, are set in the second stage of our calibration.

Housing market. The property tax rate, τ_p, is set to 1% as in Sommer et al. (2013) and the housing depreciation rate to 0.6% as estimated by Rosenthal (2014). Buying costs, τ_b, are set to 2% of the purchase price and selling costs, τ_s, are set to 6% of the sale price based on estimates by Gruber and Martin (2003). Lastly, we choose ϵ_2 so that the price elasticity of the housing supply is 0.9 as estimated by Sommer and Sullivan (2017). The level of the construction company's cost function, ϵ_1, is a free parameter that is backed out from the market clearing condition (21).

Mortgages. The interest rate premium on mortgage debt, Δ_m, is set to 1%, which is consistent with the historical difference between interest rates on 30-year government bonds and mortgage interest rates (Kaplan et al., 2019). The other mortgage parameters are set based on U.S. banking practices. We set the maximum LTV ratio, λ_1, to 0.95 so that the minimum downpayment is 5%. The maximum GDS ratio, λ_2, is set to 0.39, which is a common threshold set by lenders in the United States. The origination cost, ω_1, is set to 1% of the average household's income, which is roughly the cost of a lawyer, inspection, etc. The prepayment penalty parameter, ω_2, is set to the equivalent of 3 months' interest. The average fraction of principle paid annually, ν_1, is set to 3.33%. The prepayment penalty threshold, ν_3, is set to 18.8%. The minimum fraction of principal to be paid, ν_2, is set to 2.9%, and the threshold for full repayment, ν_4, is set to 40.3%. These last two parameters are chosen jointly to minimize the mean squared error between the payment schedules in the model and the data assuming a 30-year amortization.

4.2 Internally calibrated parameters

After assigning the parameter values listed above, there are eight parameters that still must be calibrated: the discount factor, β; the utility weight on housing services, γ; the average level of labor income taxation, τ_ℓ; the standard deduction, τ_d; the parameters of the rental company's management cost function, θ_1 and θ_2; the minimum size of a rental, $\underline{h}_r$, and the minimum size of an owner-occupied house, $\underline{h}_o$. We jointly calibrate these parameters to match eight moments in U.S. data:

- the homeownership rate, which is 65% in the 2017 wave of the AHS;
- the ratio of aggregate mortgage debt to aggregate housing wealth, which is 0.35 according to the 2019 wave of the Survey of Consumer Finances (SCF);
- the average ratio of rent to income for renters, which is 0.27 according to the AHS;
- the aggregate mortgage interest deduction as a percentage of GDP, which is 0.41% according to the Joint Committee on Taxation (JCT);
- the average labor income tax rate, which is 22.4% as estimated by McDaniel (2007);

- the rent to housing price ratio, which is 0.083 as reported by Garner and Verbrugee (2009);
- the property tax incidence of renters, which we set at 50% based on the discussion in section 2.1;
- and the fraction of renters who spend at least 50% of their gross income on rent, which is 15.1% according to the AHS.

The results of this stage of the calibration procedure are listed in panel (f) of Table 3. The internally-calibrated parameters are not individually identified, but each moment listed above plays a key role in identifying one of these parameters. The approximate mapping between moments and parameters is as follows. The mortgage debt/house value ratio pins down the discount factor, β. The lower β, the more debt homeowners are willing to take on. The average rent/income ratio determines the utility weight on housing, γ. The higher γ, the more households spend on housing relative to consumption. The average labor income tax rate identifies the labor income tax level parameter, τ_ℓ. The lower τ_ℓ, the higher the average tax rate. The MID/GDP ratio governs the standard deduction, τ_d. When τ_d rises, fewer households choose to itemize and the aggregate tax revenue lost due to the MID falls. The rent/price ratio pins down the parameter that governs the level of the rental management cost, θ_1. The higher θ_1, the higher the marginal cost of supplying rental housing, and the higher the rent-price ratio must be in equilibrium. The homeownership rate determines the minimum owned house size, $\underline{h}_o$. The higher $\underline{h}_o$, the fewer households will choose to own. Most importantly for our study, the share of property taxes borne by renters identifies θ_2, and the number of renters spending a majority of their income on housing controls $\underline{h}_r$. The identification of these parameters is discussed in detail in section 2 above.

4.3 Validation

The model performs well in reproducing other housing-related facts that we did not target in our calibration, which we report in Table 4. The values of owner-occupied houses and rental dwellings are both close to their empirical counterparts. The value of the average owner-occupied house relative to income per capita is 3.14 in the model versus 3.57 in the data, and the average value of a rental relative to the average owner-occupied house value is 59.7% in the model versus 66% in the data. The model also does well in matching statistics on mortgages. It slightly undershoots the overall fraction of homeowners who have mortgages; 51.9% of homeowners in the model have a mortgage compared to 68% in the data. However, it is quite close in terms of the number of homeowners with high-LTV mortgages; 13.6% and 9% of homeowners in the model have LTV ratios above 80% and 90%, respectively, versus 16% and 7% in the data. 33.9% of households in the model choose to take the MID instead of the standard income tax deduction, compared to 22% in the data.

Most importantly for our analysis, the model captures the fact that renters are concentrated in the bottom of the income distribution. In both model and data, renters make up most of the lowest income quintile and about half of the second income quintile, but are a minority in the top three quintiles. This means that changes in rents primarily affect low-income households. The model also reproduces closely the overall rent-to-income distribution as shown by Figure 2. We have targeted only the fraction of renters who spend at least 50% of their income on housing, but the model is consistent with the data in terms of the number of renters who spend even more than this. For example, 6% of renters spend 70% or more of their income on housing in the data, as compared to 4% in the model. Generally, slightly fewer renters spend rather large shares of their income on rent in the model as compared to the data, which indicates that the model provides a conservative assessment, if anything, of the welfare consequences of changes in rents for low-income households.

4.4 Alternative calibrations

In addition to the baseline calibration described above, we construct two alternative calibrations that represent standard approaches in the quantitative housing literature. Each alternative ignores one of the key facts described above in section 2.

In the first alternative calibration, we set $\theta_2 = 1$ so that the rental housing supply is infinitely elastic with respect to rent prices. This is the setup employed by Gervais (2002), Kaplan et al. (2019), Nakajima (2020), and Karlman et al. (2021). As we show in the appendix, it also quantitatively similar to the models of Floettoto et al. (2016) and Sommer and Sullivan (2017) in which rental housing is supplied by individual households who choose to be landlords; for example, the results reported by Floettoto et al. (2016) imply an elasticity of 38, which is 21 times larger than the elasticity in our baseline. Otherwise, this calibration is identical to the baseline in all respects, except that we recalibrate θ_1 to keep the rent-price ratio at the target value of 0.083.

In the second alternative calibration, we set $\underline{h}_r = 0$ to eliminate the lower bound on rental house size. This is the setup in Chambers et al. (2009) and Nakajima (2020). Many other studies like Floettoto et al. (2016), Sommer and Sullivan (2017), and Karlman et al. (2021) have lower bounds on rental size but do not target any moment of the rent-to-income distribution or report how this distribution looks in their models as compare to the data. Eliminating this constraint requires us to recalibrate all the other seven parameters described in section 4.2. We exclude the fraction of renters that spend more than 50% of their income on housing from the set of target moments when we do so. This calibration is similar to the baseline except in one respect: far fewer renters spend most of their income on housing. In this calibration, only 1.5% of renters spend 50% or more of their income on rent. As Figure 2 shows, this calibration clearly does a poorer job overall of matching the rent-to-income distribution than the baseline calibration, especially for higher rent-to-income ratios, which is crucial for aggregate welfare outcomes.

5 Quantitative analysis

We use our model to conduct two commonly studied experiments to illustrate the importance of the key facts discussed in section 2 for the welfare consequences of housing policy. In the first experiment, we eliminate the MID. In the second, we tax homeowners' imputed rents at the same rate as other capital income. In both, we restore fiscal balance by adjusting the standard deduction, τ_d.[10] The results of these experiments are shown in Table 5.

[10]Especially in the case of the MID, we view this assumption as the most likely political outcome. In fact, this is precisely what happened in 2018, when the government increased the standard deduction and reduced the maximum amount of mortgage debt on which one could claim the MID. In the appendix, we show our results hold also if we clear the government's budget constraint by adjusting the average labor income tax rate, τ_ℓ, instead of the standard deduction, τ_d.

Our aggregate welfare criterion is the ex-ante expected utility of a newborn household measured in consumption-equivalent terms as in Conesa et al. (2009) and Guvenen et al. (2019b). This measure is given by the formula,

$$W = \left[\frac{\int_{\mathcal{S}} V_1^{\dagger}(s) \, \mathrm{d}\Psi_1(s)}{\int_{\mathcal{S}} V_1^{*}(s) \, \mathrm{d}\Psi_1(s)} \right]^{\frac{1}{(1-\gamma)(1-\sigma)}} - 1, \tag{28}$$

where stars denote objects in the benchmark stationary equilibrium and daggers denote objects in the stationary equilibria in our policy experiments. We also compute the approval rate for each policy change as the fraction of newborns whose welfare rises,

$$A = \int_{\mathcal{S}} [V_1^{\dagger}(s) > V_1^{*}(s)] \, \mathrm{d}\Psi_1(s). \tag{29}$$

The distribution of newborns is policy-invariant because all newborns have zero financial and housing wealth, i.e., $\Psi_1^{*} = \Psi_1^{\dagger} = \Psi_1$. To dig deeper into the distributional consequences, we also compute welfare for newborns in different labor productivity groups. These distributional results are shown in Table 6.

5.1 Repealing the MID

In our first experiment, we study the effects of repealing the MID, i.e., setting $\tau_m = 0$. This policy change has been examined by numerous studies such as Gervais (2002), Chambers et al. (2009), Floettoto et al. (2016), Sommer and Sullivan (2017), Nakajima (2020), and Karlman et al. (2021), and without exception these studies have reported welfare gains in the long run.

Panel (a) of table 5 shows the main results of this change in our model, and panel (a) of Table 6 shows the distributional consequences. In our baseline model, repealing the MID would reduce aggregate welfare by 0.62% and more than 98% of newborns would experience a welfare loss. In both of our alternative calibrations, on the other hand, aggregate welfare would rise and many newborns would experience gains. This difference stems from the way the rental market responds in equilibrium and how this response affects renters, especially those at the bottom of the income distribution.

In the baseline model, repealing the MID would reduce homeownership and thus would shift demand for rentals outward. Because the supply of rentals is only moderately elastic, this would cause rents to rise by 0.61% even though house prices would fall. This would hurt all renters, but low-income households would be hit especially hard. This is because the lower bound on rental housing prevents them from downsizing, and so the increase in rents would cause a one-for-one decrease in their consumption, which is already low to begin with.

In the calibration with an infinitely elastic rental supply, homeownership would still fall, but rents would decrease by 1.04%. This is because rents do not respond to shifts in demand in this calibration, but they do respond to changes in house prices. As equation (17) shows, when the rental supply is perfectly elastic the price-rent ratio is constant, and because house prices would fall, so would rents. This would benefit all renters, especially those with low-incomes for whom the rental size constraint binds.

In the calibration without a minimum rental size, homeownership would rise slightly just as in Chambers et al. (2009), Sommer and Sullivan (2017), and Karlman et al. (2021), shifting rental demand inward. This would cause rents to fall, rather than rise as in the baseline model; here, the shift in rental demand and the decline in house prices would both push rent downward. As in the elastic-supply calibration, this would benefit low-income households, but not nearly as much as in the baseline model because they are not constrained in their choice of house size. Before moving on, it is important to highlight our finding that the prevalence of cost-burdened renters plays an important role in determining the effect of the MID on homeownership, not just welfare. This is a secondary, but still important, contribution to the literature; the effect of the MID on homeownership has been the subject of much debate as discussed by Sommer and Sullivan (2017).

5.2 Taxing imputed rents

In our second experiment, we study the effects of taxing homeowners' imputed rents. Here, we treat these rents as another form of capital income and require homeowners to pay taxes in the amount of $\tau_k(p_r - \delta)h$ as in Nakajima (2020). This policy, which has been implemented by a number of European countries,[11] has also been studied extensively (Gervais, 2002; Chambers et al., 2009; Floettoto et al., 2016; Nakajima, 2020).

As with repealing the MID, previous studies have all found long-run welfare gains from implementing this policy. Panel (b) of Table 5 shows the main results of this policy in our model, and panel (b) of Table 6 reports the distributional consequences. In our baseline calibration, taxing imputed rents would reduce aggregate welfare by 5.36% and all newborns would experience a welfare loss. In our alternative calibrations, on the other hand, aggregate welfare would rise and virtually all newborns would gain. Again, the response of the rental market in equilibrium is the source of this difference.

In the baseline calibration, taxing imputed rents would reduce homeownership substantially, leading to a large increase in demand for rentals. Although there would also be a large drop in house prices, the rental supply curve is sufficiently steep such that rents would rise by more than 6.5% in equilibrium. This is quite harmful for renters, especially those with with low incomes; newborns at the bottom of the income distribution would see welfare losses of more than 36%. Even newborns at the top of the income distribution would lose, even though they would benefit from the reduction in the standard deduction that would offset the additional tax revenue from imputed rents. This is because there is a chance, albeit a small one, of moving to the bottom of the income distribution in the future, and the large potential welfare losses associated with this switch outweigh the gains from lower taxes.

In the calibration with a perfectly elastic rental supply, homeownership and rents would both fall in this experiment as well. Again, this is because the rent-price ratio is independent of the quantity of rental housing that is supplied, so the drop in house prices would cause rents to fall. As in the previous experiment, this would benefit low-income households who are likely to be constrained by the lower bound on rental size; newborns at the bottom of the income distribution would see welfare gains of more than 17%, a complete reversal of the distributional consequences in the baseline calibration.

In the calibration without a lower bound on rental size, homeownership would fall,

[11] Imputed rents are taxed in Belgium, Iceland, Luxembourg, the Netherlands, Slovenia, Spain, and Switzerland.

which would ultimately push rents upward despite the drop in house prices. Rents would not rise as much as in the baseline model, however, and low-income households would actually gain on net. The increase in rents would have a smaller welfare bite because these households spend less on rent initially and can downsize further if necessary, and the prospect of buying a house for a lower price in the future represents a larger benefit under these circumstances.

6 Conclusion

We study how the structure of the rental market affects the consequences of homeownership subsidies. We show that two features of this structure play key roles: the price elasticity of the rental supply and the share of cost-burdened renters. The former determines how much rents respond in equilibrium to changes in demand for rental housing, and the latter determines how much renters' welfare is affected by changes in rents. When we calibrate our model to match empirical evidence on both of these features, it predicts that removing the MID and taxing imputed rents would reduce aggregate welfare and would particularly hurt low-income renters. When we omit either of these pieces of evidence from our calibration, the model predicts that these policies would increase welfare.

Our results stand in sharp contrast to previous findings in the literature. Numerous studies, such as Gervais (2002), Chambers et al. (2009), Floettoto et al. (2016), Sommer and Sullivan (2017), Nakajima (2020), and Karlman et al. (2021), have found that these same policies would be welfare-improving. Our analysis shows that this is because all of these studies fail to account for at least one of the key features of the rental market that we have emphasized. Except for Chambers et al. (2009), all have rental supply elasticities that are far higher than the evidence indicates, and only Gervais (2002) has a realistically high share of cost-burdened renters. We show that accounting for both of these features is crucial.

Our findings are particularly important for major metropolitan areas where homeownership rates are lower and many more renters are cost-burdened as compared to the national average.[12] In New York City, for example, the homeownership rate is only 32%

[12] Oates and Fischel (2016) argue that the supply of rental housing is also less elastic in urban areas.

and 26% of renters are severely cost-burdened (NYU Furman Center, n.d.). In these areas, removing homeownership subsidies is likely to be substantially more harmful, both in the aggregate and for low-income renters in particular. Further, rents nationwide have grown more quickly than incomes over the past two decades, leading to a large increase in the share of cost-burdened renters (Whitney, 2023). If this trend continues, eliminating homeownership subsidies would be even more harmful in the future. More broadly, our findings have additional implications for the effects of many other policy changes on housing affordability. For example, The Federal Reserve has raised interest rates dramatically to combat inflation in the wake of the COVID-19 pandemic, and mortgage debt has become significantly more costly to take on as a result. Our findings suggest that this could put significant upward pressure on rents and materially hurt low-income households.

References

Arslan, Yavuz, Bulent Guler, and Temel Taskin (2015) 'Joint dynamics of house prices and foreclosures.' *Journal of Money, Credit and Banking* 47(1), 133–169

Barnett, C. Lance (1979) 'Expected and Actual Effects of Housing Allowances on Housing Prices.' *Real Estate Economics* 7(3), 277–297

Carroll, Robert, and John Yinger (1994) 'Is the property tax a benefit tax? the case of rental housing.' *National Tax Journal* 47(2), 295–316

Chambers, Matthew, Carlos Garriga, and Don E. Schlagenhauf (2009) 'Housing policy and the progressivity of income taxation.' *Journal of Monetary Economics* 56(8), 1116–1134

Chatterjee, Satyajit, and Burcu Eyigungor (2015) 'A quantitative analysis of the u.s. housing and mortgage markets and the foreclosure crisis.' *Review of Economic Dynamics* 18(2), 165–184

Collinson, Robert, and Peter Ganong (2018) 'How do changes in housing voucher design affect rent and neighborhood quality?' *American Economic Journal: Economic Policy* 10, 62–89

Conesa, Juan Carlos, S. Kitao, and Dirk Krueger (2009) 'Taxing capital? not a bad idea after all.' *American Economic Review* 99, 25–48

Díaz, Antonia, and María José Luengo-Prado (2008) 'On the user cost and homeownership.' *Review of Economic Dynamics* 11(3), 584–613

England, Richard W. (2016) 'Tax incidence and rental housing: A survey and critique of research.' *National Tax Journal* 69(2), 435–460

Fack, Gabrial (2006) 'Are housing benefit an effective way to redistribute income? evidence from a natural experiment in france.' *Labour Economics* 13(6), 747–771

Floettoto, Max, Michael Kriker, and Johannes Stroebel (2016) 'Government intervention in the housing market: Who wins and who loses?' *Journal of Monetary Economics* 80, 106–123

Garner, Thesia, and Randal Verbrugee (2009) 'Reconciling user costs and rental equivalence: Evidence from the us consumer expenditure survey.' *Journal of Housing Economics* 18(3), 172–192

Gervais, Martin (2002) 'Housing taxation and capital accumulation.' *Journal of Monetary Economics* 49, 1461–1489

Gibbons, Stephen, and Alan Manning (2006) 'The incidence of uk housing benefit: Evidence from the 1990s reforms.' *Journal of Public Economics* 90, 799–822

Gruber, Joseph, and Robert Martin (2003) 'Precautionary savings and the wealth distribution with illiquid durables.' Discussion paper 773, Federal Reserve Board International Finance

Guvenen, Fatih, Fatih Karahan, Serdar Ozkan, and Jae Song (2019a) 'What do data on millions of u.s. workers say about life cycle labor income risk?' Working paper

Guvenen, Fatih, Gueorgui Kambourov, Burhannetin Kuruscu, Sergio Ocampo, and Daphne Chen (2019b) 'Use it or lose it: Efficiency gains from wealth taxation.' Working paper

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni Violante (2017) 'Optimal tax progressivity: An analytical framework.' *Quarterly Journal of Economics* 132(4), 521–539

Herbert, Chris, Alexander Hermann, and Daniel McCue (2018) 'Measuring housing affordability: Assessing the 30 percent of income standard.' Working paper, Joint Center for Housing Studies, Harvard University

Kangasharju, Aki (2010) 'Housing allowance and the rent of low-income households.' *Scandinavian Journal of Economics* 112, 595–617

Kaplan, Greg, Kurt Mitman, and Giovanni Violante (2019) 'The housing boom and bust: Model meets evidence.' *Journal of Political Economy, Forthcoming*

Karlman, Markus, Karin Kinneruda, and Kasper Kragh-Sørensena (2021) 'Costly reversals of bad policies: The case of the mortgage interest deduction.' *Review of Economic Dynamics*

Larrimore, Jeff, and Jenny Schuetz (2017) 'Assessing the severity of rent burden on low-income families.' https://doi.org/10.17016/2380-7172.2111. Accessed: 2023-12-02

McDaniel, Cara (2007) 'Average tax rates on consumption, investment, labor, and capital in the oecd 1950–2003.' Working paper

Nakajima, Makoto (2020) 'Capital income taxation with housing.' Journal of Economic Dynamics and Control

NYU Furman Center 'State of the city 2021: State of renters and their homes.' https://furmancenter.org/stateofthecity/view/state-of-renters-and-their-homes. Accessed: 2023-12-02

Oates, Wallace E., and William A. Fischel (2016) 'Are local property taxes regressive, progressive, or what?' *National Tax Journal* 69(2), 415–433

Orr, Larry (1968) 'The incidence of differential property taxes on urban housing.' *National Tax Journal* 21, 253–262

Orr, Larry L. (1970) 'The incidence of differential property taxes: A response.' *National Tax Journal* 23, 99–101

Rosenthal, Stuart (2014) 'Are private markets and filtering a variable source of low-income housing? estimates from a "repeat income" model.' *American Economic Review* 104, 687–706

Rotberg, Shahar (2022) 'The implications of housing for the design of wealth taxes.' *International Economic Review*

Rotberg, Shahar, and J. B. Steinberg (2022) 'Tax Evasion and Capital Taxation.' Working paper

Sayag, Doron, and Noam Zusman (2020) 'Who benefits from rental assistance? evidence from a natural experiment.' *Regional Science and Urban Economics*

Schwegman, David, and John Yinger (2020) 'The shifting of the property tax on urban renters: Evidence from new york state's homestead tax option.' Working paper

Sommer, Kamila, and Paul Sullivan (2017) 'Implications of us tax policy for house prices, rents, and homeownership.' *American Economic Review* 108(2), 241–274

Sommer, Kamila, Paul Sullivan, and Randal Verbrugge (2013) 'The equilibrium effect of fundamentals on house prices and rents.' *Journal of Monetary Economics* 60(7), 854–870

Susin, Scott (2002) 'Rent vouchers and the price of low-income housing.' *Journal of Public Economics* 83, 109–152

Viren, Matti (2013) 'Is the housing allowance shifted to rental prices?' *Empirical Economics* 44, 1497–1518

Whitney, Peyton (2023) 'Number of renters burdened by housing costs reached a record high in 2021.' https://www.jchs.harvard.edu/blog/number-renters-burdened-housing-costs-reached-record-high-2021. Accessed: 2023-12-02

Wiehe, Meg, Aiden Davis, Carl Davis, Matt Gardner, Lisa Christensen Lee, and Dylan Grundman (2018) 'Who pays: A distributional analysis of the tax systems in all 50 states.' 6th edition, The Institute on Taxation and Economic Policy

Table 1: List of quantitative studies of homeownership policies

Study	Rental supply elasticity	Min. rental size target	Repealing MID	Taxing imputed rents
Gervais (2002)	Infinite	Rent/spending of lowest inc. quintile	Gain	Gain
Chambers et al. (2009)	Endog. landlords	No min. size	Gain	Gain
Floettoto et al. (2016)	Endog. landlords	No min. size	Gain	LR gain, SR loss
Sommer and Sullivan (2017)	Endog. landlords	Not reported	Gain	–
Nakajima (2020)	Infinite	No min. size	Gain	Gain
Karlman et al. (2021)	Infinite	Not reported	LR gain, SR loss	–

Table 2: List of empirical evidence on the rental supply elasticity

Study	Estimate	Geography
(a) Renters' property tax incidence		
Carroll and Yinger (1994)	11%	Boston, U.S.
Schwegman and Yinger (2020)	14%	New York City, U.S.
Fack (2006)	22%	Paris, France
Orr (1968)	30%	Boston, U.S.
Orr (1970)	46%	Boston, U.S.
Wiehe et al. (2018)	50%	Nationwide, U.S.
(b) Share of vouchers captured by landlords		
Sayag and Zusman (2020)	25%	Jerusalem, Israel
Viren (2013)	33%	Nationwide, Finland
Gibbons and Manning (2006)	40%	Nationwide, U.K.
Collinson and Ganong (2018)	46%	Nationwide, U.S.
Kangasharju (2010)	65%	Nationwide, Finland
Susin (2002)	100%	Nationwide, U.S.

Table 3: Calibration

Parameter	Description	Value	Source
(a) Demographics and preferences (assigned)			
J	Lifespan	60	Live from 26-85
J_R	Retirement age	41	Retirement at age 66
σ	Risk aversion	2	Standard
$\{\xi\}_{j=1}^{J}$	Equivalence scale	Varies	AHS (2017)
(b) Labor productivity process (assigned)			
σ_z	Dispersion in labor market productivity	0.202	Rotberg and Steinberg (2022)
ρ_z	Persistence of labor market productivity	0.937	Rotberg and Steinberg (2022)
$\{\zeta_j\}_{j=1}^{R}$	Life-cycle labor market prod.	$1+\min\{\frac{0.38(j-1)}{30}, 0.38\}$	Guvenen et al. (2019a)
$\kappa(z)$	Social security benefit	Varies	Guvenen et al. (2019b)
(c) Taxes and interest rate (assigned)			
r	Interest rate	3%	Kaplan et al. (2019)
τ_k	Capital gains tax rate	15%	Karlman et al. (2021)
τ_e	Personal exemption	0.43× std. deduction	U.S. tax code
ψ	Progressivity of labor income taxes	0.85	Heathcote et al. (2017)
(d) Housing market (assigned)			
τ_p	Property tax rate	1%	Sommer and Sullivan (2017)
δ	Depreciation rate	0.6%	Rosenthal (2014)
τ_b	Buying cost	2%	Gruber and Martin (2003)
τ_s	Selling cost	6%	Gruber and Martin (2003)
ϵ_2	Price elasticity of housing investment	1.9	Sommer and Sullivan (2017)
(e) Mortgages (assigned)			
Δ_m	Mortgage interest rate premium	1%	Kaplan et al. (2019)
λ_1	Max. LTV ratio	95%	US banking regulations
λ_2	Max. GDS ratio	39%	US banking regulations
ω_1	Origination cost	1% of avg. income	See text
ω_2	Prepayment penalty	3 months' interest	US banking practices
υ_1	Avg. fraction of debt paid	3.33%	Avg. % of debt paid annually
υ_2	Min. fraction of debt paid	2.9%	30-year amortization
υ_3	Threshold for prepayment penalty	18.8%	US banking practices
υ_4	Threshold for full repayment	40.3%	Pct. of home value paid last period
(f) Jointly calibrated			
β	Discount factor	0.938	Mortgage debt/housing wealth=35%
γ	Utility weight on housing preferences	0.161	Total rent/total income=27%
τ_ℓ	Avg. labor income tax	0.850	Avg. labor income tax rate=22.4%
τ_d	Standard deduction	4% of avg. labor inc.	MID/GDP=0.41%
θ_1	Rental management cost level	$\theta_1 S^{\theta_2-1}/p_r$=29.8%	Rent/housing price=8.3%
θ_2	Rental management cost convexity	2.95	Renters' tax incidence=50%
$\underline{h_r}$	Smallest rental	0.290× avg. owned house size	Renters with rent/income≥ 0.5=15.1%
$\underline{h_o}$	Smallest owned house	0.435× avg. owned house size	Homeownership rate=65%

Table 4: Non-targeted moments

Statistic	Baseline	Data	Source
Avg. owned house value (× income per HH)	3.14	3.57	SCF (2019)
Avg. rental value (% avg. owned house value)	59.7	66	Chatterjee and Eyigungor (2015)
Homeowners with a mortgage (%)	51.9	68	SCF (2019)
Homeowners with LTV≥80 (%)	13.6	16	SCF (2019)
Homeowners with LTV≥90 (%)	9	7	SCF (2019)
Households who take the MID (%)	33.9	22	JTC (2010)
Renters by income quintile (%)			
First	78.73	61.37	AHS (2017)
Second	51.15	47.52	AHS (2017)
Third	28.33	38.16	AHS (2017)
Fourth	12.07	28.18	AHS (2017)
Fifth	2.73	15.41	AHS (2017)

Table 5: Aggregate effects of housing policy changes

Calibration/ Experiment	House price (% chg.)	Rent (% chg.)	HO rate (p.p. chg.)	Welfare (% chg.)	Approval (%)
(a) Repealing MID					
Baseline model	-0.81	0.61	-0.69	-0.62	1.27
Elastic supply	-0.93	-0.50	-1.04	0.26	41.73
No min. rental	-1.92	-0.67	0.10	0.12	57.76
(b) Taxing imputed rents					
Baseline model	-5.71	6.64	-6.38	-5.36	0.00
Elastic supply	-6.97	-3.75	-12.35	3.70	100.00
No min. rental	-7.08	1.25	-4.15	1.38	98.73

Table 6: Welfare effects across the income distribution

z (% of pop.)	z_1 (2.26)	z_2 (9.51)	z_3 (22.87)	z_4 (30.69)	z_5 (22.87)	z_6 (9.51)	z_7 (2.26)
(a) Repealing MID							
Baseline	-2.93	-0.79	-0.23	-0.59	-0.64	-0.01	0.34
Elastic supply	2.45	0.79	0.48	-0.16	-0.54	0.08	0.41
No min rental	0.26	0.29	0.29	-0.06	0.04	-0.03	0.22
(b) Taxing imputed rents							
Baseline	-36.77	-9.20	-1.97	-1.10	-1.28	-0.28	-0.04
Elastic supply	17.66	6.10	4.43	2.02	0.16	0.45	0.47
No min. rental	0.59	1.16	2.24	1.36	0.60	0.19	-0.03

Figure 1: Illustration of rental supply elasticity identification

(a) Sensitivity of rents to demand

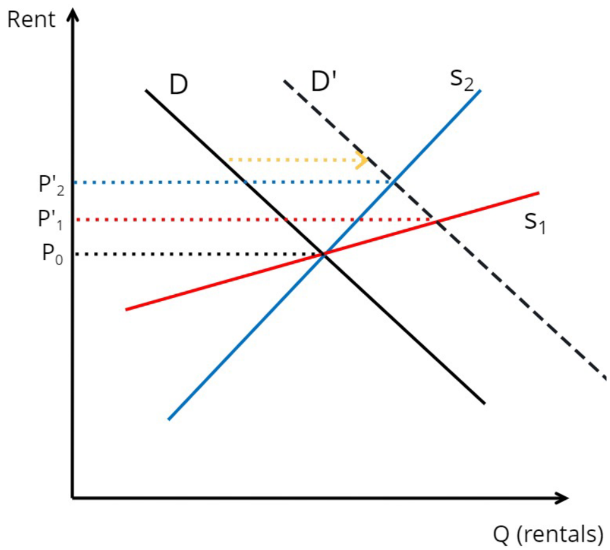

(b) Identification from property tax incidence

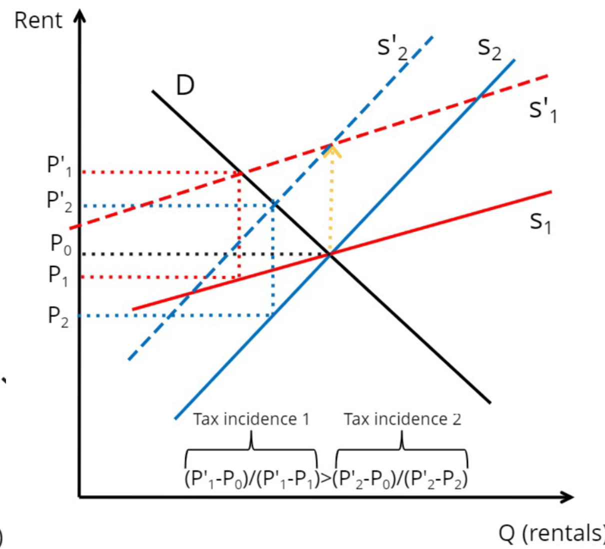

Figure 2: CDF of the rent to income ratio

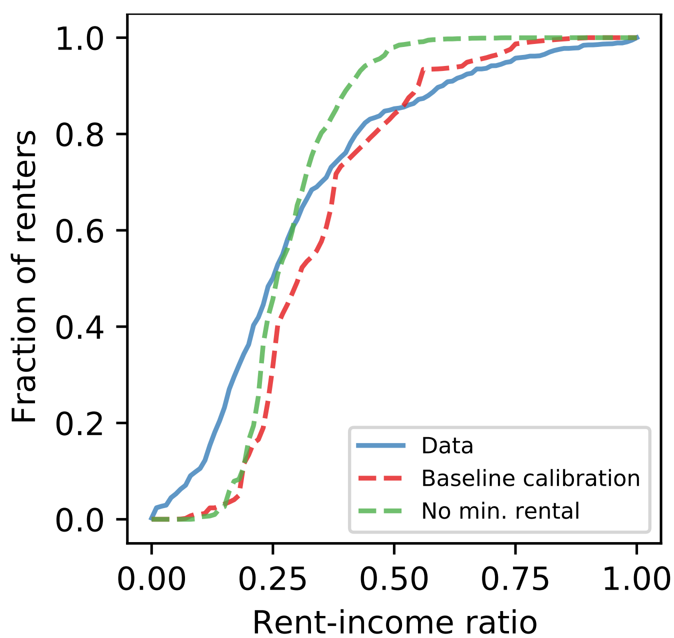

Appendix (For Online Publication)

Appendix A describes data sources and the steps taken to calculate calibration moments using these data. Appendix B describes additional sensitivity analyses we have performed.

A Data

Most of our assigned parameter values and calibration targets are taken directly from the literature. There are a few, however, that we constructed ourselves from several data sources.

A.1 2017 American Housing Survey

The 2017 American Housing Survey (AHS) is an important data source for us because it provides the information we need to construct the distribution of rent-to-income ratios shown in Figure 2. We use the file `household.csv` in the national public-use data files.[13] We include all observations for renters for whom the household head's age is between 26 and 85 (inclusive) and rent (variable: RENT) to income (variable: HINCP) ratio is less than 1. We measure the rent-to-income ratio as the former divided by the latter. We also compute the number of renters in each income quintile using the categorical variable TENURE.

A.2 2019 Survey of Consumer Finances

We also use the 2019 wave of the Survey of Consumer Finances (SCF) to measure several of the non-targeted moments we use to validate our model.[14] We measure the average owned house value using the variable HOUSES. We measure the number of homeowners with a mortgage by computing the number of homeowners with MRTHEL>0. We measure the number of homeowners with LTV ratios above 0.8 and 0.9 as those with MRTHEL/HOUSES>0.8 and MRTHEL/HOUSES>0.9, respectively.

[13] See https://www.census.gov/programs-surveys/ahs/data.2017.List_1739896299.html#list-tab-List_1739896299.

[14] See https://www.federalreserve.gov/econres/scfindex.htm.

A.3 Long-term mortgage payment schedules

We obtain a mortgage principal payment schedule from https://www.calculator.net/mortgage-calculator.html, using the parameter values listed in the text. Our calibration strategy is to minimize the mean squared error between this schedule and the payment schedule implied by the parameters ν_1 and ν_4. Specifically, we choose ν_1 to minimize this MSE and then set ν_4 to ensure that the principal is completely paid off in year 30. Note that although the calibrated value of ν_4 may appear large, under the actual payment schedule 41.5% of the principal is paid off in the last 5 years of the mortgage.

B Sensitivity analysis

Here we consider two sensitivity analyses that demonstrate the robustness of our findings. In the first, we study experiments in which the government's budget is cleared by adjusting the average labor income tax rate parameter, τ_ℓ, instead of by adjusting the standard deduction. In the second, we study an alternative model of rental supply in which homeowners can choose to become landlords as in Chambers et al. (2009), Floettoto et al. (2016), and Sommer and Sullivan (2017).

B.1 Adjusting income tax rates instead of the standard deduction

In the main text of the paper, we clear the government's budget constraint (23) in our policy experiments by adjusting the standard deduction, τ_d. Here, we study experiments in which we adjust τ_ℓ, which controls the average labor income tax rate, instead. Part (1) of table B.1 shows the results of these experiments. All the results are quite similar to the main results in table 5. Most importantly, removing the MID and taxing imputed rents would still reduce welfare in the baseline calibration, but would raise welfare in the high-elasticity and no-min-rental calibrations.

B.2 Model with endogenous landlords

In our baseline model, rental housing is supplied by a representative management corporation. This is the same approach taken by several other papers that have studied homeownership subsidies, most notably Gervais (2002), Nakajima (2020), and Karlman et al. (2021). This model has two attractive features for our purposes: (i) one can easily control the rental supply elasticity by changing the convexity of the rental management cost; and (ii) this elasticity can be directly identified from the share of property taxes borne by renters. However, there is another popular model in which rentals are supplied by individual households who choose to be landlords. This approach is used by Chambers et al. (2009), Floettoto et al. (2016), and Sommer and Sullivan (2017). Here, we show that this "endogenous landlord" model has a high rental supply elasticity and that the results of our policy experiments in this model are similar to the results from the elastic-supply calibration of our baseline model.

B.2.1 Model formulation

In the endogenous-landlord model, rental housing is supplied by homeowners who choose to rent out a portion of their houses. Specifically, a household that owns h' units of housing can choose to rent out $x \leq h' - \underline{h}_r$ units; just like renters, landlords must consume at least $\underline{h}_r$ units of housing services. As in the other models in this literature, landlords must pay a fixed cost of ϕ and earn rental income of $p_r x$. Rental income is taxed at the capital income tax rate τ_k. The fixed cost, depreciation, property taxes and mortgage interest on the rental property are allowed to be deducted. Thus, taxable rental income is $p_r x - \phi - px(\delta + \tau_p) - m'(x/h')r$. We allow this taxable income to enter the GDS constraint, which means that landlords can obtain bigger mortgages than non-landlords. Only the portion of mortgage interest that is attributable to owner-occupied housing, $m'(1 - x'/h')r$, can be deducted from labor income taxes. The dynamic program in this model is

$$V_j(s) = \max_{c,a',o',h',n,m',x} \left\{ u_j(c, h' - x) + \beta \int_{\mathcal{Z}} V_{j+1}(s') \, \mathrm{d}F(z, z') \right\} \tag{B.1}$$

subject to

$$\begin{aligned} c + a' + r_m m + (1 - o')p^r h + o \left[\delta + \tau_p + \mathbb{1}_{\{o'=0 \vee h' \neq h\}} \tau_s \right] ph + o' \mathbb{1}_{\{o=0 \vee h' \neq h\}} (1 + \tau_b) ph' \\ = y_j(z) - \tau_j(z, o, h, m) + [1 + r(1 - \tau_k)]a + m' - m - n\omega_1 - (1 - n)\tau_{pp}(m, m') \\ + p_r x - \mathbb{1}_{\{x>0\}}\phi - \tau_k[p_r x - \phi - px(\delta + \tau_p) - m'(x/h')r] \end{aligned} \tag{B.2}$$

$$a' \geq 0, \tag{B.3}$$

$$h' \geq o' \underline{h}_o + (1 - o')\underline{h}_r \tag{B.4}$$

$$n \in \{0, o' \mathbb{1}_{\{j<j_R\}}\} \tag{B.5}$$

$$x \in [0, o' \max(0, h' - \underline{h}_r)] \tag{B.6}$$

$$nm' \leq \lambda_1 ph' \tag{B.7}$$

$$n[\tau_p ph' + (r_m + \nu_1)m'] \leq \lambda_2[y_j(z) + ra + p_r x - \phi - px(\delta + \tau_p) - m'(x/h')r] \tag{B.8}$$

$$(1 - n)m' \leq \mathbb{1}_{\{m \geq \nu_4 ph\}}(1 - \nu_2)m \tag{B.9}$$

B.2.2 Calibration and rental supply elasticity

We calibrate this model to match the same set of target moments as before with two changes. First, the fixed landlord cost, ϕ, is chosen during our internal calibration procedure so that 10% of homeowners are landlords as estimated by Chambers et al. (2009). Second, and most important for our purposes, we no longer target the property tax incidence of renters because there is no parameter in this model that controls the aggregate rental supply elasticity. When we measure the aggregate supply elasticity in our calibrated endogenous-landlord model, however, we find an extremely high number. Near the calibrated benchmark equilibrium, the aggregate rental supply curve is extremely shallow: holding other prices fixed, a 1% increase in the rental price leads to a 31% increase in the quantity of rental housing supplied. This implies that renters' property tax incidence is 91%, far higher than the estimates in the literature reported in table 2 and essentially the same in concept as the infinitely elastic supply case.

The supply elasticity in our endogenous-landlord model is very similar to the elasticity implied by the results of Floettoto et al. (2016). In their experiment in which they eliminate the MID, rents rise by 2% and the stock of rental housing supplied in equilibrium rises by 76%, implying an aggregate rental supply elasticity of 38.[15] Most other papers that use the endogenous-landlord model like Sommer and Sullivan (2017) do not report enough information to make this calculation. The one exception is Chambers et al. (2009), who report that a 4% increase in rents increases supply by 6.3% in partial equilibrium, which implies a rental supply elasticity of only 1.6. We have tried many configurations of the endogenous-landlord model and cannot produce an elasticity anywhere near this low number. We can reduce the elasticity of the rental supply curve to 7 by taxing rental income at the same progressive rate as labor income, but this implies a property tax incidence for renters of 85% which is still far higher than in the data.[16] As we discuss below, our endogenous-landlord model generates results that are very similar to those of Sommer and Sullivan (2017) and Floettoto et al. (2016), especially when we have no minimum rental size.

We have also attempted to build a version of the endogenous-landlord model with a lower aggregate supply elasticity generated by a convex rental management cost as in our baseline model. However, because (i) each individual landlord supplies small amounts of rental housing relative to the aggregate quantity supplied, and (ii) aggregate supply changes are driven by the extensive margin as well as the intensive margin, we have found that it is not possible to identify the landlord-level management cost convexity from the aggregate property tax incidence. We speculate that one needs to model a management cost that depends on the aggregate quantity supplied, not the individual, landlord-level quantity. For example, one might be able to generate a low aggregate rental supply elasticity if they modelled a rental management company that landlords pay to manage their properties, and this company has increasing marginal costs. This

[15] In their experiment, the aggregate stock of housing, the house prices, and the landlordship rate do not change much, implying that the rental supply curve does not shift much. However, if anything it is likely that the rental supply curve shifts inward because the homeownership rate falls. If so, this would mean that the supply elasticity in their model is even higher than we have calculated.

[16] Moreover, in this version of the model the rental supply increases dramatically more in our policy experiments. The net effect is that rents fall even more than in the infinite-elasticity corporate model, leading to even larger welfare gains.

would create an externality in the rental market, however, because individual landlords would not take into account how their choices affect the aggregate management cost. Grappling with this issue is outside the scope of our paper.

B.2.3 Results

Part (2) of Table B.1 shows the results of our policy experiments in the endogenous-landlord model. We conduct these experiments in two versions of the model. In the first (labeled "large min. rental"), we calibrate the minimum rental size to match the share of cost-burdened renters as in the baseline calibration. In the second (labeled "no min. rental") we set the minimum rental house size to zero.

In the large-min-rental version of the model, removing the MID and taxing imputed rents would reduce both house prices and rents, just as in the elastic-supply calibration of our baseline model, although the movements in prices are somewhat larger. Consequently, aggregate welfare would rise substantially in both experiments. Thus, this sensitivity analysis shows that the endogenous-landlord model has a highly elastic rental supply curve, and this is precisely why papers that use this model find welfare gains from removing the MID and taxing imputed rents. However, this version of the endogenous-landlord model is not directly comparable to other papers in the literature that do not match the share of cost-burdened renters.

In the no-min-rental version of the endogenous-landlord model, the results are qualitatively the same, but the movements in both prices and welfare are smaller. The results of the MID-removal experiment in this version of the model are in all respects very similar to those of Sommer and Sullivan (2017). This model predicts that removing the MID would increase homeownership, just as that paper found, and raise aggregate welfare by 0.83%, which is close to that paper's 0.76%. Although Sommer and Sullivan (2017) have a minimum rental house size that is allowed to be different from the minimum owned house size, they do not report the value of the former. Our results here strongly suggest, however, that in Sommer and Sullivan (2017)'s benchmark there are far fewer cost-burdened renters than in the data.

Table B.1: Aggregate effects of housing policy changes: sensitivity analysis

Calibration/ Experiment	House price (% chg.)	Rent (% chg.)	HO rate (p.p. chg.)	Welfare (% chg.)	Approval (%)
(1) Using τ_ℓ instead of τ_d					
(a) Repealing MID					
Baseline model	-0.77	0.75	-0.88	-0.86	8.99
Elastic supply	-0.86	-0.46	-1.33	0.07	41.73
No min. rental	-1.91	-0.69	0.06	0.04	65.48
(b) Taxing imputed rents					
Baseline model	-4.96	4.46	-5.28	-3.24	8.99
Elastic supply	-5.72	-3.08	-9.00	2.90	100.00
No min. rental	-6.41	1.16	-4.07	0.86	100.00
(2) Endogenous landlords					
(a) Repealing MID					
Large min. rental	-3.63	-4.46	-0.04	3.34	100.00
No min. rental	-2.19	-2.68	0.15	0.83	100.00
(b) Taxing imputed rents					
Large min. rental	-10.10	-11.19	-12.54	8.23	100.00
No min. rental	-7.02	-9.31	-12.67	3.32	100.00