CPB Netherlands Bureau for
Economic Policy Analysis

Mortgage Debt Limits and Buy-to-Let Investors: A Structural Model of Housing with an Endogenous Rental Sector

Home ownership rates have decreased since the financial crisis. We show that part of this development can be explained by decreased access to mortgage credit. As households are able to borrow less to purchase a home, their demand shifts to the rental sector. Buy-to-let investors respond to this demand by converting houses into rentals. This mechanism can explain about a fifth of the increased number of rental houses in the market-rate sector in the Netherlands.

Taxing rents or imposing a cap on the number of rentals can partially counteract the effects of mortgage debt limits on the number of rentals. However, such measures also push households that prefer to rent into owner-occupation and can thus be counterproductive.

CPB Discussion Paper

Jurre Thiel, Henrik Zaunbrecher

September 2023

Doi: https://doi.org/10.34932/wfaj-zs98

Mortgage Debt Limits and Buy-to-Let Investors: A Structural Model of Housing with an Endogenous Rental Sector*

Jurre H. Thiel[†], Henrik Zaunbrecher[‡]

September 22, 2023

Abstract

We investigate the effects of mortgage debt limits on household tenure choice and housing market equilibrium. Based on a natural experiment that changed Debt-Service-to-Income (DSTI) limits in the Netherlands, we show that a 1% decrease in the maximum mortgage leads to a 8% increase in the probability to rent rather than own a house. To understand the equilibrium effects of DSTI limits, we build a structural model of the Dutch housing market. The model features an endogenous rental sector in the form of buy-to-let investors that convert owner-occupied housing into rentals. Our results show that DSTI limits cause houses to shift into the rental sector. This causes a deadweight loss as households with a preference for owner-occupation rent in equilibrium. Policymakers must therefore balance the macroprudential benefits of mortgage debt limits against their potential distortive effects on the housing market. Finally, we show that taxing buy-to-let investors, a common policy to promote home ownership, can only partially revert this deadweight loss.

1 Introduction

Since the Great Financial Crisis, many advanced economies have experienced a decline in home ownership rates.[1] This inevitably means that more households rent. Meanwhile, buy-to-let investors, who turn homes into rentals, have become more prominent actors on the housing market.[2] This has sometimes led to push back from policy makers against

*We are grateful to Emile Cammeraat, Hans Koster, Jante Parlevliet and Jan Rouwendal for helpful comments and suggestions. Results based on calculations by CPB Netherlands Bureau for Economic Policy Analysis using microdata from Statistics Netherlands.

†CPB Netherlands Bureau for Economic Policy Analysis. j.h.thiel@cpb.nl.

‡CPB Netherlands Bureau for Economic Policy Analysis. h.w.zaunbrecher@cpb.nl.

[1]For example, in the US, the home ownership rate dropped from a peak of 69% to 63% in 2016, recovering slightly since to 66% in 2022. In the UK, the home ownership rate decreased from 73% in 2007 to 63% in 2016.

[2]For example, in some American metropolitan areas, over 10% of housing market transactions were undertaken by buy-to-let investors between 2014 and 2016 (Mills et al., 2019).

buy-to-let investing. For example, the UK has increased the stamp duty on buy-to-let investors.[3]

In this article we focus on a simple explanation for the increasing share of rentals: constrained access to mortgage credit. As a reaction to high debt levels that ultimately contributed to the financial crisis, macroprudential policies have limited household access to mortgage across the board (Claessens, 2015). There is an intuitive link between access to mortgage credit and tenure choice, as households that cannot borrow enough to buy their preferred house may choose to rent it instead. There is also empirical evidence, going back at least to Linneman (1985), that documents the effects of mortgage constraints on tenure choice. However, we do not know to which extent tightened debt limits can explain the increase in rentals versus other explanations. Nor do we know who gains and loses when debt limits are tightened and when policies that deter buy-to-let investing can be justified.

We look at these issues in the context of the Dutch housing market, where the market-rate rental market has tripled in size between 2006 and 2018 (Figure 1a).[4] As in most advanced economies, there was a tightening of mortgage lending standards in the same period. For the average household, the Debt-Service-To-Income (DSTI) limit, which measures what percentage of its income a household may spend on its mortgage, has decreased by 17% (Figure 1a).[5] This has also led to policy concerns about accessibility of the housing market. In 2019, 23.6% of renters were prevented by DSTI limits from purchasing a similar home (Figure 1b). Among renting households in the lowest three income deciles of our sample, 57.1% were not able to purchase a similar home.[6]

To start, we exploit a natural experiment to show that mortgage credit constraints influence households' tenure. Between 2015 and 2019, DSTI limits were gradually relaxed for couples but not for single households. This increased couples' average maximum mortgage by €3,344 or 1.95% compared to singles. Using a triple differences strategy, we show that this change caused couples that were renting and constrained by DSTI limits in 2015 to be 15.6% less likely to rent in 2019. This results in an elasticity of probability to rent with respect to mortgage debt limits of -8.

To compute the welfare effects of DSTI limits in the presence of buy-to-let investors, we then build a structural model of the Dutch housing market. We also use the model to evaluate policies, like a rent tax, that are often used by policy makers to deter buy-to-let investing. The major empirical challenge that we face is that our model must be able to explain why some households rent and some own a home. Since Poterba (1984) and Case and Shiller (1989), it is common to assume that households' tenure choices follow from a financial no-arbitrage condition in which the user cost of owning a home is compared to the rent. However, as stressed by Glaeser and Gyourko (2007), this approach is limited for several reasons. First, the costs of owning a home depend on factors such as expected

[3] `https://www.gov.uk/stamp-duty-land-tax/residential-property-rates`. Accessed May 17, 2023.

[4] We focus on market-rate rentals, excluding the social rental market. See Section 2 for more details.

[5] In The Netherlands, the legal DSTI limit automatically decreases with the mortgage interest rate. Hence, we compute the decrease holding the interest rate constant, so that this measures a decrease in DSTI limits due to policy changes only.

[6] As we exclude households that live in the social housing sector, these households are not in the poorest three deciles of the full population. The average household income of these households in 2019 is €35,332, which is approximately the median income across all households.

Figure 1: Developments in the Dutch rental market

(a) Number of rentals and mortgage credit limits (b) Renters by income

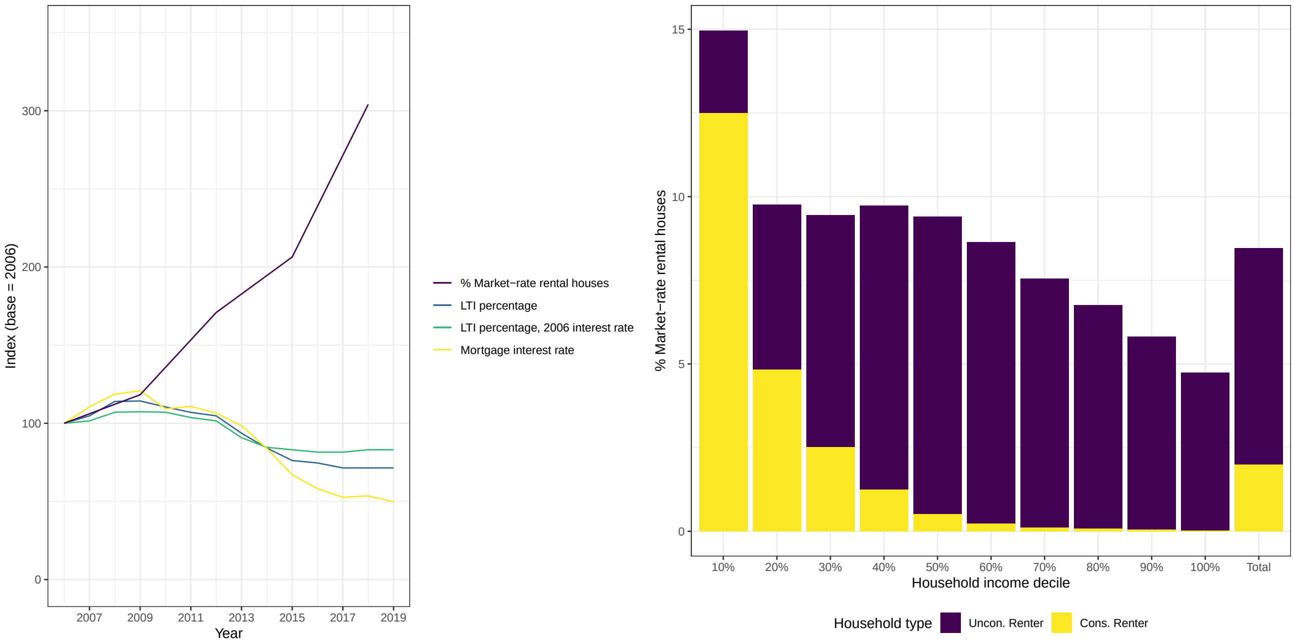

Note: The left panel shows the fraction of rental houses in the market-rate sector, i.e. excluding social housing, over time. The base year is 2006. The number of rental of houses is based on the WooN-survey by Statistics Netherlands, which is held every 3 years. Hence, observations are only available for 2006, 2009, 2013, 2015 and 2018. The DSTI percentages are computed by the authors based on the average household income in each year. The mortgage interest rate is the average interest rate on 5-10 year fixed rate mortgages, 10-year mortgages being the most common mortgages in the Netherlands, as computed by the Dutch central bank. The right panel shows the fraction of household active in the market-rate housing sector that rents a house in 2019. Constrained renters are constrained by DSTI limits in that they cannot buy a similar house as they rent. This is computed based on our consolidated data set (see Section 2). We define a "similar" house as house of the same housing type, where housing types are computed using K-means clustering (see Section 4.1). Income deciles are with respect to our sample, i.e. households that live in a market-rate house.

maintenance and appreciation that are difficult to observe. Hence, it is difficult to say when a household "should" prefer to own a home from a financial standpoint. Second, rental homes are different from owner-occupied homes and renters are different from owners.[7] This complicates a purely financial approach since it assumes that households choose between directly comparable homes. Moreover, as indicated by Figure 1b, there is likely heterogeneity in tenure preferences: households that rent but could buy clearly prefer renting, while restricted renters could well prefer to own.

To overcome these conceptual issues, we estimate household tenure preferences directly from the data. To do so, we set up a dynamic model of tenure choice in the vein of Poterba (1984), which we extend with heterogeneous housing. Households compare the discounted utility of buying a house with that of renting. We show that two parameters form a sufficient statistic for tenure preferences: the derivatives of utility with respect to *current* house prices and rents. These derivatives can be obtained from standard static discrete choice models. In this way, we are able to simplify a dynamic problem into a simpler static one.

[7] For example, Glaeser and Gyourko (2007) find for the US that rental homes are more likely to be apartments, while owner-occupied homes are more likely to be single-family dwellings. They also find that owners are on average older and richer than renters.

In particular, we extend Bayer et al. (2004, 2007) to incorporate both rental and owner-occupied properties. In their model, households make a discrete choice amongst houses.[8] We include rental and owner-occupied homes separately in the choice set and allow households to have different preferences over them. Households are constrained by DSTI limits if they want to purchase a house. We show that the user cost of housing, which measures a household's tenure preference, can be computed directly from the model's primitives. Attractive for our purposes is that, in the model, both houses and households are heterogeneous. The model can thus rationalize why some household groups are more likely to rent than others.

On the supply side, our model features heterogeneous buy-to-let investors. The number of rentals is determined by a no-arbitrage condition, which says that, in every segment of the housing market, the rent-price ratio must equal the marginal investor's hurdle rate of capital. We show that the rental supply curve can be identified non-parametrically based on this condition. The main idea is that for houses with a higher rent-price ratio, investors make a larger return and hence supply more rentals. Because the rent-price ratio is likely correlated with supply shocks, we develop an instrument that shifts the demand for the rentals. The instrument is based on the intuition that preferences for house characteristics (e.g. size) are correlated with tenure preferences. Since these preferences follow from our demand model, we can compute an instrument that shifts the demand to rent a particular type of house.

Turning to our results, we find that the tightening of DSTI limits between 2013 and 2019 can explain 21% of the increase in rentals in this time frame. Hence, DSTI limits distort households' tenure choice and, as a consequence, the equilibrium number of rentals. However, this also shows that there also is a role for other explanations, such as changes to lending standards, the low-interest rate environment or changes to the taxation of owner-occupied housing (Lambie-Hanson et al., 2022). Changes to DSTI limits change the number of constrained renters. For example, a uniform decrease of DSTI limits of 10% causes an increase of constrained renters of 1.12 percentage points (11.4%) amongst households with the lowest 30% incomes in our sample.[9]

The deadweight loss of DSTI limits in the housing market can be significant. For example, we find that a relatively small increase of debt limits of 5% from 2019 levels would increase total surplus in the housing market by €1.3 billion per year, equivalent to a 0.2% increase in gdp. This gain is realized because an increase in borrowing capacity unlocks two avenues for gains from trade. The first, which we dub the tenure choice channel, follows from current renters with a large preference for home ownership being able to purchase their home. The second, which we call the house choice channel, consists of an increase in allocative efficiency from households that were already able to buy being able to move into previously unavailable houses. Since, as we noted above, the effect on the absolute number of (constrained) renters is relatively small, we conclude that the house choice channel is the more important determinant of the deadweight loss of DSTI

[8]The original formulation also includes household sorting across neighbourhoods. As this is not the main focus of our analysis, we omit this feature of the model.

[9]We focus on this sub-sample as most constrained renters occur within this sub-sample of the population (Figure 1b), but this result generalizes to the full sample.

limits. However, the effects of DSTI limits are non-linear: the opposite policy, tightening DSTI limits by 5%, has almost no effects on total surplus. Policy makers should weigh deadweight losses introduced through distortions of tenure and house choice against positive effects of DSTI limits on financial stability and paternalistic motives to protect households from excess debt.

While relaxing DSTI limits would increase total surplus, they do not make households without property ("outsiders") better off. This result follows directly from our model assumption that the housing stock is fixed. When the supply of housing is inelastic, changes in demand due to expansions of DSTI limits are capitalized into house prices. Changes to total surplus hence always accrue to property owners in our model. The assumption that the housing supply is fully inelastic is clearly strong. However, housing supply elasticities in The Netherlands are typically estimated to be low (Michielsen et al., 2017). It is therefore likely that changes to DSTI limits only lead to small changes in the surplus of outsiders even in a model with a more realistic supply side.

We also use our structural framework to evaluate policies that deter buy-to-let investing. Such measures may be taken when policy makers want to help households that prefer to buy but are restricted to the rental sector, but still value the macroprudential and/or paternalistic benefits of DSTI limits. In particular, we consider an increased tax on rent and a cap on the percentage of rentals. Policies of this kind have been introduced in the Netherlands as well as abroad. For example, in the UK, buy-to-let properties are subject to higher capital gains taxes and stamp duty. In some Canadian jurisdictions, further foreign purchases of houses are banned, which functions as a soft cap on the number of buy-to-let properties.

We find that an (increased) rent tax can counteract the deadweight loss caused by DSTI limits, but to a limited extent. A rent tax reduces the equilibrium fraction of rentals as supplying rentals becomes less profitable. In this way, it allows some constrained households with a preference for ownership to buy a house. A rent tax hence reduces the deadweight loss of DSTI limits through the tenure choice channel but not the house choice channel: the choice sets of existing homeowners are not expanded. Moreover, we find that the size of the tenure choice channel is highly heterogeneous across housing market segments. As a result, a uniform tax rate cannot eliminate the distortion for some segments of the housing market without also reducing the proportion of rentals in others.

However, the most beneficial tax rate we consider, 2.5%, only raises total surplus with €381 million in total. The effect is thus about 30% of that of a small relaxation of DSTI limits. A tax, however, has two additional benefits over DSTI limits. The first is that a rent tax does not impose negative externalities on the financial system. Second, a rent tax can be used to make housing market outsiders better off by redistributing its revenues toward them.

In the counterfactuals we consider, a cap on the amount of rentals always lower total surplus. We hence provide suggestive evidence that taxing rentals is a less costly than restricting the number of rentals to increase the accessibility of the buy-to-let market. A reason for this is that, in addition to the mechanisms we described for the rent tax, a quantity cap introduces waterbed effects: the number of rentals increases for segments of the housing market that are not directly affected by the cap.

Related literature. This paper contributes to a small but growing literature on the housing market effects of macroprudential policies. We extend this literature in two ways. First, it has largely focused on the effects of Loan-To-Value (LTV) limits, which limit the loan principal relative to the value of the house (Bekkum et al., 2019; Tzur-Ilan, 2023). This literature broadly finds support for the idea that macroprudential policies influence the housing market, as well as tenure choices more specifically. Closest to our work is Bekkum et al. (2019), who show that in The Netherlands LTV limits influence households' tenure choice by making constrained household less likely to buy a house. We show that macroprudential policies that limit the principal with respect to income, such as DSTI and Loan-To-Income limits, also have large economic effects, in particular on tenure choice. Second, our structural model allows us to move beyond reduced form estimates and quantify the equilibrium effects of macroprudential policies. This enables us to also measure the effects of macroprudential policies on households that are not directly constrained by them.

To do so, we build on a literature in urban economics on tenure choice. As we described above, it has been common to assume tenure choices follow from a financial comparison, where the rent is compared against the user cost of housing (Poterba, 1984; Linneman, 1985; Hill and Syed, 2016). In our model, households make a similar comparison. However, we extend this literature by estimating households' user cost instead of calibrating it. We also show there is significant heterogeneity in the user cost of housing.

Finally, our work also relates to a literature on the role of buy-to-let investors on housing markets. Perhaps the closest paper to ours is Rouwendal et al. (2023), who consider the relationship between DSTI limits and buy-to-let investors in a theoretical model. Their theoretical arguments support our empirical results that DSTI limits encourage buy-to-let investors to convert owner-occupied houses into rentals. Our results differ in that they find no distortive effects of DSTI limits on the matching of houses to households in the presence of buy-to-let investors. They obtain this result because they assume that the supply of rentals by buy-to-let investors is fully elastic. Therefore, an investor is always willing to buy a house from and rent it back to a household as soon it becomes capital constrained. We relax this assumption by estimating the supply curve of rentals.

Another structural model of a housing market with buy-to-let investors is (Han et al., 2022), who use it to estimate the effects of a transaction tax (stamp duty). In their model, households exogenously rent or own, while in our model the tenure choice is endogenous. Our model also features heterogeneity in both houses and households. This allows us to disentangle tenure preferences from preferences for housing characteristics and to compute redistributive effects. But, like other macroeconomic models with a rental sector (Kiyotaki et al., 2011; Sommer et al., 2013; Favilukis et al., 2017), the model in Han et al. (2022) is a dynamic macroeconomic model. They can hence speak to general equilibrium effects, for example on the interest rate, and transition dynamics. We view our approach as complementary to the macroeconomic approach.

The remainder of this article is structured as follows. In Section 2, we describe the Dutch housing market and our data set. Section 3 provides reduced-form evidence on the effects of DSTI limits on tenure choice. We develop our structural model in Section 4, while we present our main results in Section 5. Section 6 concludes.

2 The Dutch housing market

The Dutch housing market is divided into two segments: market-rate and social housing. Social housing comprised 37% of the housing stock in 2018 (Statistics Netherlands, 2022). These houses are rent-controlled, with the maximum rent determined by the government based on property characteristics.[10] Social housing is typically only available to households that are low-income or otherwise vulnerable and is typically heavily rationed.[11]

In the market-rate sector, owner-occupied properties are dominant: only 12% of the market-rate sector comprised rental properties in 2019 (Figure 1b). The small size of the market-rate rental sector is typically attributed to the tax subsidies owner-occupiers receive. These advantages include the deductibility of mortgage interest payments and an exemption of house equity in capital gains taxation. The net fiscal subsidy is estimated to be in the 20-25% range (Ewijk and Lejour, 2017).

Even though the amount of market-rate rental homes is low, its number has increased substantially, tripling between 2006 and 2018 and increasing by 70% between 2013 and 2018 (Figure 1a). As in other markets where the rental sector has grown since the Great Financial Crisis, there are likely multiple reasons for this growth (Lambie-Hanson et al., 2022): more stringent lending standards, households being unable to move because they have negative home equity, a large inventory of homes for investors to buy due to depressed demand from households, and an inflow of capital into buy-to-let investment due to low interest rates. In the Dutch context, tax subsidies for owner-occupation were also reduced after the crisis, which likely also contributed to an increased fraction of rentals. Our goal is not to quantify all possible causes of the increased investor activity in the Dutch housing market. Rather, we focus on the first cause: more stringent lending standards.

Mortgage lending standards in the Netherlands primarily take the shape of Debt-Service-To-Income (DSTI) limits.[12] A DSTI limit maximizes a household's total spending on its mortgage, i.e. the sum of interest and mortgage payments, as a percentage of its income. This percentage is determined yearly by the government and itself depends on income, i.e. higher incomes can spend a larger fraction of their income on debt service. Contrary to most other countries, Loan-to-Value (LTV) limits, that maximize the loan principal to a multiple of income, are not strict in the Netherlands. Until 2013, the maximum LTV ratio was 105%, after which it was gradually decreased to 100% in 2018. Since households can still borrow the full value of the house, the DSTI limit is hence much more likely to

[10]Whether a house is social and hence rent-controlled is determined by a government-set points system. Points are given for house characteristics such the square footage. A house is rent-controlled if and only if the number of points is below a cut-off. If a house is rent-controlled, the landlord may charge a fixed rent per point. The maximum rent in the social sector was €720.42 per month in 2019.

[11]As the income limit is only binding at the moment the household moves in, it is possible for higher-income households to live in social housing if their incomes increased afterwards.

[12]Sometimes, these are called Payment-to-Income (PTI) limits. In the Dutch policy debate, the term Loan-to-Income (LTI) limit is often used. This is technically not correct, as the limit is on the total mortgage spend and not the size of the loan, although there is obviously a relationship between the two.

bind.[13] As Figure 1a shows, DSTI limits were also reduced more than LTV limits during the same period: controlling for changes in the mortgage interest rate, the average household's maximum mortgage expenditure decreased with 17% between 2013 and 2018. We therefore restrict our analysis to the effect of DSTI limits.

To analyze the effects of DSTI limits, we build a consolidated dataset of the Dutch housing market based on rich microdata provided by Statistics Netherlands.[14] Our data set comprises the period 2015-2019. While earlier data are available, many significant reforms to the Dutch mortgage market were undertaken in 2013 that could potentially confound our estimates.[15] Our data contain all households and dwellings in the Netherlands. As our focus is on the interaction between the rental and owner-occupied market, we exclude social housing, which contains only rentals. We also have a pragmatic argument for this exclusion: our structural model is based on a revealed preferences approach, in which households live in their most desired house. As social houses are rationed and can have long wait lists, this assumption is not valid for this segment. We similarly exclude households that live in social housing, as the the market-rate sector is typically too expensive them. We do include households that are on the margin of the market-rate housing sector, but currently live in social housing or live with another household.[16] We do this so that our structural model includes the extensive margin of the market-rate housing market. Furthermore, we consider only homes with a single household and exclude e.g. nursing homes, some forms of student accommodation and or jails.[17]

While we have rich register data on most parts of the housing market, this is lacking for some aspects of the rental sector. In particular, while transaction prices are available for all owner-occupied properties, we have to rely on survey data for the rental sector. We confront this issue by imputing rent using a machine-learning algorithm, gradient boosted trees (Ke et al., 2017). We first impute whether a rental is social or market-rate, and then impute the rent for market-rate rentals only. Imputations are based on household and house characteristics. We provide more details in Supplementary Material B, where we also show our imputations have an excellent fit in a hold-out sample.

While we have data on almost 5 million unique dwellings in our data set, it is too computationally demanding to model the demand for every individual house. However, assuming housing is homogeneous or can be described by a single quality index is limiting: rentals are different from owner-occupied houses, as we in Table 1 and previously stressed by Glaeser and Gyourko (2007). Owner-occupied houses tend be more more expensive, larger and located in less urban areas. They are also significantly more likely to be single-family homes rather than apartments. Renters tend to be younger and poorer than owner-occupiers and comprise smaller households.

[13]This does not imply that a 100% LTV limit is not constraining at all. Households cannot borrow to pay the transaction costs related to purchasing a house, such as stamp duty or notarial costs. These costs typically amount to 3 – 5% of the price of the house, such that some savings are required.

[14]An overview of the data sets we use is provided in Supplementary Material A

[15]For example, the deductibility of mortgage interest payments was restricted to amortizing mortgages.

[16]We do this by including households that move into the market-rate housing sector in year t in our data set at year $t-1$. The idea is that these households are most likely to move into the market-rate sector if prices change.

[17]In these cases it is unknown how the house characteristics are split across households (e.g. how large the individual rooms/flats for each household are), and thus we can't use the revealed preference approach to infer their housing preferences.

Table 1: Summary statistics of houses and occupants

	All houses	Owner-occupied	Rentals
Houses			
Government-assessed value (€)	320136.8	321807.2	303305.1
	(190489.4)	(189452.1)	(199864.5)
Size (m^2)	138.4	140.9	113.0
	(86.3)	(84.8)	(96.0)
Urbanity (higher = more urban)	3.3	3.2	4.0
	(1.3)	(1.3)	(1.2)
% apartments	19	15	60
Households			
Household size	2.5	2.5	2.0
	(1.3)	(1.3)	(1.1)
Average age all adults	54.1	54.8	47.0
	(15.9)	(15.4)	(18.6)
Gross household income (€/ year)	98065.4	100075.4	77938.5
	(150871.1)	(150431.5)	(153763.0)
N	4575843	4152638	423205

Note: The table shows average characterstics of houses and households for 2019. Standard errors are in parentheses.

We hence compromise by dividing the houses in our data set into a smaller number of *house types* using a K-means clustering algorithm. This machine learning method automatically combines observations into clusters that have similar observable characteristics. It does this by choosing K centroid points in characteristics space such that the total squared distance between all observations and the closest centroid is minimized. We cluster all houses into $K = 5,865$ types based on the government-assessed value ("WOZ" in Dutch), square footage and their location as measured by their coordinates.[18] We stress that we do *not* cluster on whether a house is a rental or owner-occupied. Hence, most housing types are available in both sub-markets, but some housing types will have more rentals than others.

To make this procedure computationally achievable on our large data set, we use mini-batched K-means (Sculley, 2010). This algorithm speeds the clustering process up by iteratively batching on a random subset of the data. Further details of our clustering procedure are provided in Supplementary Material C. Here we also show that our house types are also similar for variables we do not explicitly cluster on, such as transaction prices and year of construction.

[18] This number is based on the largest number of centroids that we could estimate our model for, which was 6,000. After clustering, a number of housing types was dropped from the analysis as we had some clusters that were empty in some years.

Figure 2: Maximum mortgage for single and joint borrowers, 2015–2019

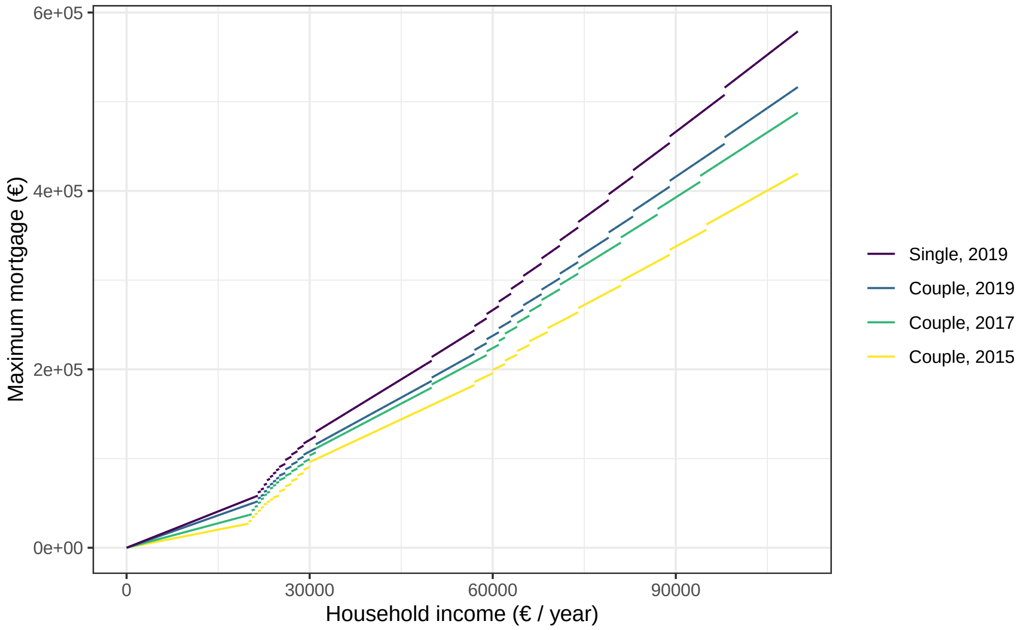

Note: The figure shows the maximum mortgage allowed under DSTI limits as a function of household income, for single and joint borrowers (couples). For couples, it is assumed that the lowest income is 36% of household income. This was the average in 2016, and is based on research on the financial situation single- and double-income households in the Netherlands from Statistics Netherlands (2018). For single households only the maximum mortgage in 2019 is shown as there are no significant changes compared to 2015.

3 The effects of a debt-limit reform

We start by providing reduced-form evidence on the effect of mortgage debt limits on households' tenure choice. We do so by exploiting a reform in the Dutch lending regulations that increased DSTI limits of joint borrowers relative to single borrowers. In the Dutch system, the maximum DSTI ratio is a function of income, with higher-income households permitted to spend a larger fraction of their income on debt service. For joint borrowers, the lowest income is discounted when determining this maximum ratio. Up to and in 2015, joint borrowers' lowest income was multiplied with a factor of 33%. This means that a couple with incomes €60,000 and €30,000 were allowed to spend the same percentage of their household income as a single borrower with an income of €70,000. This ratio was increased to 50% in 2016, 60% in 2017 and 70% in 2018.[19] As a result, joint borrowers' lending capacity increased. As Figure 2 shows, this increased the borrowing capacity of joint borrowers compared to single borrowers with a similar household income.[20]

To estimate the causal effect of this expansion of borrowing capacity on the propensity to rent, we focus on directly-affected households. In particular, we look at households that are renting in 2015 and are restricted by DSTI limits from purchasing a similar home ("restricted households"). A similar home is a home of the same house type, as we defined in Section 2. We then compare couples with single households to estimate the effect on

[19]It was further increased after our sample ends, until it reached 100% in 2023.

[20]To compute the maximum borrowing capacity, we obtain the maximum debt service for a given income level from the official government tables. We then compute the largest a household can obtain, for which the debt service is below this maximum. Here, we assume a 30-year annuity mortgage and take average mortgage interest rate of new originations in the given year.

tenure choice.

There are possible spillover effects between treated and untreated restricted households due to equilibrium effects of the reform on rents, house prices and the number of rental homes. Therefore, we additionally control for the difference in the effect between unrestricted couples and singles. This leads to the following triple difference specification:

$$y_{it} = \beta_0 + \beta_1 \text{Couple}_i + \beta_2 \text{Constrained}_i + \sum_s \gamma_s T_s + \beta_3 \text{Couple}_i \times \text{Constrained}_i + \sum_s \delta_s \text{Couple}_i \times T_s + \sum_s \zeta_s \text{Constrained}_i \times T_s + \sum_s \theta_s \text{Couple}_i \times \text{Constrained}_i \times T_s + \varepsilon_{it}. \quad (1)$$

Here, y_{it} is a dummy that equals one when household i rents in year t. $Couple_i$ indicates whether household i is a couple and hence is treated by the debt limit reform. $Constrained_i$ measures whether the household was a constrained renter in 2015, and hence directly affected by the reforms. T_s contains year dummies and $(\beta, \gamma, \delta, \zeta, \theta)$ are coefficients to be estimated. The coefficients of interest, θ_t, measure the average treatment effect of the debt limit expansion on the treated. In our case, these are couples that are restricted renters in 2015.

Our identifying assumption is that, absent the expansion of debt limits, the relative probability of couples to rent relative to singles trends in the same way for restricted and unrestricted households (Olden and Møen, 2022). In practice, this assumption is threatened because restricted couples show a larger propensity to have significant income growth than unrestricted couples. As a result, their lending capacity increases independently of any changes to debt limits. We control for this in various ways. First, we exclude households that have a large income increase in our sample period.[21] Second, we control for household income and wealth. To gauge the robustness of our estimates, we consider two specifications for the "untreated, unaffected" group. We use all households whose tenure choice is not restricted by DSTI limits (including homeowners), as well as a subsample where we only consider unrestricted renters. In the latter case, our sample only contains households that rent in 2015. The idea is that if our approach to control for differences in income growth between restricted and unrestricted couples is valid, our estimates should be robust to the exact specification of restricted couples.

Figure 3 shows the average propensities to rent for the groups in our triple-differences setup. It shows a clear divergence in rental probabilities between restricted couples and singles as borrowing constraints of the former become more relaxed relative to the latter. Some single borrowers that were restricted in 2015 did also purchase a home in the following years. This can happen because their financial situation improved, so that they are no longer restricted, or because they chose to purchase a house that is of lower quality than the house they rented in 2015.

We find a significant and economically meaningful impact of borrowing constraints on

[21] To be precise, we estimate the ratio of the maximum and minimum household incomes for the years 2015–2019. We restrict the sample to households for whom this ratio is below the 85th percentile. To test whether this removes diverging trends in income and wealth, we re-estimate equation (1) with household income and wealth as dependent variables. As required, we find no effect of the expansion of debt limits on income and wealth after this restriction.

Figure 3: Average rental probabilities for single and joint borrowers

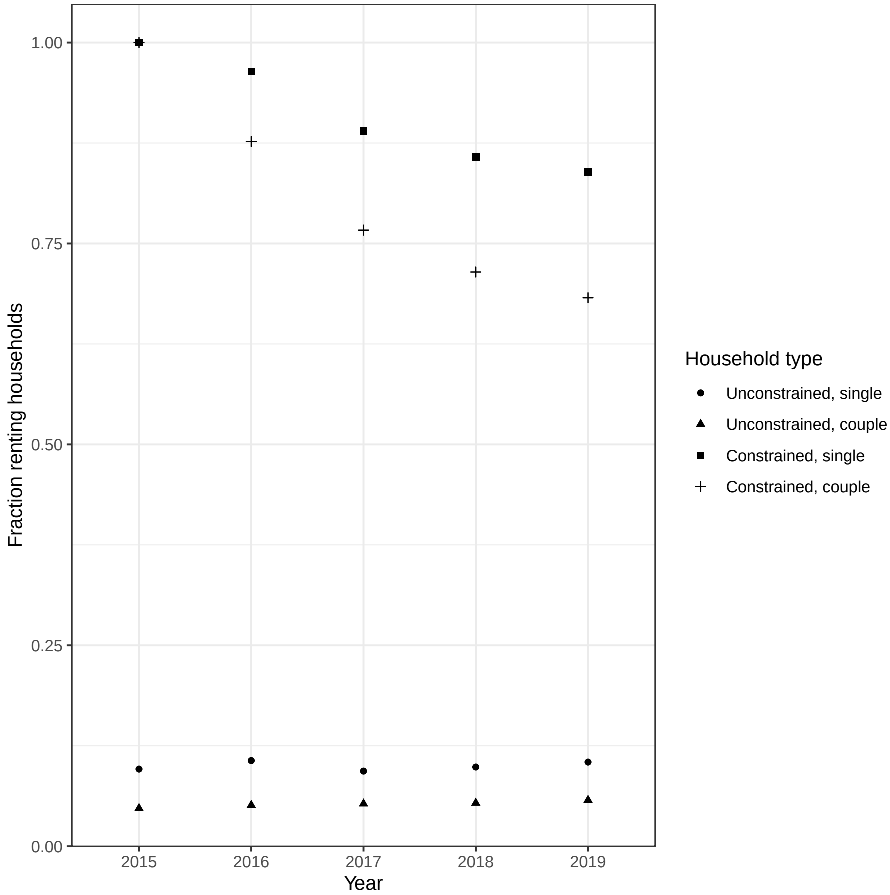

Note: The figure show the effect of the increases of DSTI limits for co-borrowers between 2015–2019. It shows the average probability to rent for restricted, unrestricted, single and couple households. Restricted households are defined as households that rent in 2015 and are restricted by DSTI limits to purchase a similar home.

tenure choice across our specifications (Table 2). Our preferred specification is one in which we *both* exclude households with large income changes and control for income and wealth. This specification is the only one for which the effects on the probability to rent and the maximum mortgage is robust to constraining our sample to contain only renters (columns 3 and 6). In this specification, couples who were constrained in 2015 are 4 percentage points (15.6%) less likely to rent compared to singles. Because their maximum mortgage debt increased with €3344 (1.95%), this means that the elasticity of rental probabilities with respect to debt limits is -8. This estimate pertains to households that are both treated and restricted, i.e. it measures the average treatment effect on the treated. For unrestricted households, the elasticity should be zero. Since about 3% of households are constrained renters, a back-of-the-envelope calculation then suggests that the average elasticity, i.e. the average treatment effect, is -.21.

Table 2: The effects of a debt limit expansion on the probability to rent

	(1)	(2)	(3)	(4)	(5)	(6)
Probability to rent						
2019 × Treat	-0.038	-0.154	-0.038	-0.040	-0.107	-0.040
	(0.000)	(0.001)	(0.000)	(0.000)	(0.004)	(0.001)
Maximum mortgage						
2019 × Treat	1457.472	3156.372	3299.714	11617.126	5664.800	3244.113
	(0.000)	(424.912)	(77.516)	(0.000)	(1296.451)	(814.848)
Elasticity	-17.225	-27.567	-7.500	-2.209	-11.285	-8.082
Exclude large income changes	Yes	No	Yes	Yes	No	Yes
Control for income & wealth	No	Yes	Yes	No	Yes	Yes
Only renters	No	No	No	Yes	Yes	Yes
Couple × Constrained FE	Yes	Yes	Yes	Yes	Yes	Yes
Couple × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Constrained × Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	17002691	19948155	17002691	941387	1486969	941387

Note: The table shows the causal effect of the 2015-2019 debt limit expansion for joint borrowers compared to single borrowers on two outcomes of interest: the probability to rent and the maximum allowed mortgage according to DTI limits. The effect is estimated through the triple differences model (1). From these estimates and (unreported) baseline values, we compute the elasticity of renting with respect to debt limits. When we exclude large income changes, we drop all households whose maximum income between 2015 and 2019 was more than 2.25 times the minimum period. In the "Only renters" arm, we exclude owner-occupiers from the untreated, unaffected group, so that it only contains unconstrained renters. Standard errors are in parentheses.

4 A structural model of a housing market with buy-to-let investors

Our reduced-form results show that households are sensitive to debt limits in their tenure choice. This analysis, however, is not sufficient to compute the deadweight loss of debt limits. First, as we have not estimated the supply of rental houses, we are unable to compute the equilibrium price effects of changes to debt limits. Second, because debt limits not only distort tenure choice, but potentially also house choice. That is, even households that choose to own may own a lower quality house than they would under more relaxed debt limits. The distortion of tenure and house choice likely interact. For example, a small relaxation of debt limits can cause a household to move from a rental into a lower-quality owner-occupied house. In this case, tenure choice becomes less distorted whereas the distortion of house choice increases. We need a model of the housing market that considers both distortions to accurately compute welfare effects.

We therefore build and estimate a structural model of the Dutch housing market. Apart from computing the deadweight loss of debt limits, the model also allows us two perform two exercises that are relevant to policy makers. First, we are able to compute distributive effects and see, for example, whether housing market outsiders are affected disproportionately. Second, policy makers are typically reluctant to relax debt limits but can nevertheless be concerned about the distortive effects of debt limits on the housing market. We can use our structural model to compute to what extent other policies, such as taxing rentals, can address these distortions.

Our model contains four key ingredients: i) it contains both a rental and owner-occupied market; ii) the tenure status of houses is endogenous; iii) households are constrained by DSTI limits in the owner-occupied market; iv) because the extent to which DSTI limits bind differs between households and housing market segments, both households and homes are heterogeneous.

4.1 Environment

Tenure choice. We extend the tenure choice model of Poterba (1984) with heterogeneous houses. Household i chooses on house type j and tenure $o \in \{R, O\}$, where R indicates the house is rented and O that it is owner-occupied. We derive demand for a single household. We assume all parameters are household-specific, but we omit household subscripts for readability. Our focus for now is to derive a tractable model of tenure choice. We therefore postpone the inclusion of some salient details, such as debt limits, to below.

The model is dynamic and households are assumed to have infinite lives. A house of type j provides flow utility v_{jt}. Every year, a household quits its current tenure with the exogenous probability κ. This can be interpreted as a change in preferences, for example due to significant life events as getting children or finding a new job. The household has a discount rate ρ.

For rentals, the initial rent is p^R_{jt} and real growth rate of rents is ψ. When choosing to buy, the household does so at a price p^O_{jt}, for which it makes a down-payment θp^O_{jt}. The real interest rate is r. Interest payments are subsidized at an effective rate χ.[22] House prices grow at a real rate g. We assume all payments are made at the end of the year.

The utility of renting house type j is then

$$u^R_{jt} = \sum_{s=1}^{\infty} \left(\frac{1-\kappa}{1+\rho}\right)^s \frac{1}{1-\kappa} \left(v^R_{jt} - \alpha(1+\psi)^{s-1} p^R_{jt}\right) = \frac{v^R_j}{\rho+\kappa} - \frac{\alpha}{\rho+\kappa+\kappa\psi-\psi} p^R_{jt}, \quad (2)$$

where α is the household's marginal utility of income. The first term is the net present value of the flow utility obtained by living in rental house j, where the second term subtracts the net present value of rent payments in utility terms.

The utility of buying a house of type j is

$$u^O_{jt} = -\alpha\theta p^O_{jt} + \sum_{s=1}^{\infty} \left(\frac{1-\kappa}{1+\rho}\right)^s \frac{1}{1-\kappa} \left\{v^O_j + \alpha\left(-r(1-\chi)(1-\theta)p^O_{jt} + \kappa((1+g)^{s-1} - (1-\theta))p^O_{jt}\right)\right\}.$$

The first term is the down payment on the house. In the sum, the first term again is the flow utility of housing consumption. The second term contains the interest payment on the mortgage, where the term $(1-\chi)$ corrects for any tax advantages on mortgage interest. The final term is the value of appreciation at the time of sale. The latter two terms are converted to utility terms through multiplication with the marginal utility of income α.

[22]In The Netherlands, the major subsidies for owner-occupation are the tax deductibility of mortgage interest payments and the exclusion of equity returns of owner-occupied dwellings from income taxation (Ewijk and Lejour, 2017).

This discounted utility can be written as

$$u^O_{jt} = \frac{v^O_{jt}}{\rho + \kappa} - \alpha \left(\frac{1 - (1 - \rho - \kappa)\theta + r(1 - \chi)(1 - \theta)}{\rho + \kappa} - \frac{\kappa}{\rho + \kappa + \kappa g - g} \right) p^O_{jt}. \tag{3}$$

Since utility is ordinal, we can multiply (2) and (3) with $\rho + \kappa$ to obtain the following static discrete choice formulation for household i:

$$\begin{aligned} u^R_{ijt} &= v^R_{ijt} - \alpha^R_{it} p^R_{jt}, \\ u^O_{ijt} &= v^O_{ijt} - \alpha^O_{it} p^O_{jt}. \end{aligned}$$

House and tenure choice depend on the utility of the different house types and by tenure preferences. The tenure preferences are encapsulated by the coefficients α^R_{it} and α^O_{it}, which are household-specific functions of the deep parameters $(\alpha, \rho, \kappa, \psi, g, \theta, \xi)$. For example, a household that is more likely to move (higher κ) will have different tenure preferences, as will a household that expects more growth of house prices (higher g).

In the case where owner-occupied and rental houses provide the same flow utility, i.e. $v^R_{ijt} = v^O_{ijt}$, we obtain the original model by Poterba (1984). The household then only chooses to rent or to own. It owns when $u^O_{ijt} \geq u^R_{ijt}$, or when

$$\frac{p^R_{jt}}{p^O_{jt}} \geq \omega_{it} \equiv \frac{\alpha^O_{it}}{\alpha^R_{it}}.$$

As households buy instead of rent whenever the rent-price ratio exceeds the user cost of housing (Poterba, 1984), the ratio of price coefficients equals the user of cost housing. We denote the user cost with ω_{it}.

Demand. To obtain an estimable model, we parameterize the flow utilities v^o_{ijt} as a function of observable house type characteristics and unobserved factors. In particular, we write

$$v^o_{ijt} = X'_{jt}\beta_{it} - \xi^o_{jt} + \varepsilon^o_{ijt}.$$

Here, X_{jt} contains observable (to us) house characteristics, ξ^o_{jt} is a market-level measure of unobserved house quality and ε^o_{ijt} is an idiosyncratic error term. With this formulation, we obtain a discrete choice model for housing in the vein of Bayer et al. (2004):

$$\begin{aligned} u^R_{ijt} &= X'_{jt}\beta_{it} - \alpha^R_{it} p^R_{jt} + \xi^R_{jt} + \varepsilon^R_{ijt}, \\ u^O_{ijt} &= X'_{jt}\beta_{it} - \alpha^O_{it} p^O_{jt} + \xi^O_{jt} + \varepsilon^O_{ijt}. \end{aligned} \tag{4}$$

Hence, while tenure choice is a dynamic problem, we can estimate it as if it were a static choice. The reason is that the price coefficients equal the user cost of housing, which completely determines tenure preferences. The ratio of price coefficients thus forms a sufficient statistic for tenure preferences.

Additionally, we add the option for a household to live in the outside option, which gives utility

$$u_{i0t} = \varepsilon_{i0t}.$$

A household choosing the outside option does not live in the market-rate housing sector. Hence, it lives in the social housing sector or with another household such as family.

Household i is constrained by DSTI limits. That is, its choice sets contains all owner-occupied houses that it is allowed to purchase given its maximum mortgage and its wealth, including any equity it has in its current house. We observe a small number of households owning houses that they would not be allowed to purchase under current DSTI limits, which may reflect that they purchased their house when DSTI limits were more relaxed. We assume that a household can always buy any owner-occupied house that it owns in the data.[23] Households can choose to live in any rental house.[24] Hence, we denote by $\mathcal{C}_{it}$ the houses that household i can *buy*.

Observable characteristics X_{jt} contribute equally to the utility of owner-occupied houses and rentals, but the role of prices p_{jt} and unobservable characteristics ξ_{jt} differ. In our specification, a house type j has two associated prices: for rentals, the yearly rent p^R_{jt}, for owner-occupied houses, the purchase price p^O_{jt}. Since these are in different scales, they have different coefficients. The unobserved quality can also differ, to reflect that even observably similar rentals can differ from owner-occupied dwellings in unobserved ways. For example, rental houses might be less-well maintained because owner-occupiers have different maintenance incentives than renters (Henderson and Ioannides, 1983; Halket et al., 2020). As a result of this divergence between the utility of a rental and an owner-occupied house, a household typically prefers one ownership mode over the other for the same house type.

An advantage of this approach is that it allows us to estimate, rather than assume, households' tenure preferences. As we showed above, the user cost of housing ω_{it} can be computed as the ratio of the price coefficients on rental and owner-occupied houses, which are directly estimable. Moreover, since we allow the price coefficients to be heterogeneous, so is the user cost of housing. The alternative would be to take a stand on the determinants of the user cost of housing. However, these determinants can be difficult to observe (in the case of discount rates or expectations) and are likely heterogeneous (Glaeser et al., 2013).

However, the consequence is that the user cost is fixed in our model. When we look at the determinants we have used to derive the user cost of housing, we can see that most are indeed exogenous. However, the growth rates of house prices and rents, g and ψ, as well as moving probabilities, κ, are not. However, since the effects on price *levels* are minor in all our counterfactuals, typically below 5%, we do not expect large endogenous feedback loops on the growth rate of prices. Similarly, moving probabilities can change with the supply of rentals, as renters are typically more likely to move than owner-occupiers (Head and Lloyd-Ellis, 2012). However, we find relatively small effects on the supply of rentals. We thus also believe that the overall effect on moving probabilities would be small in our

[23]Without this assumption, the likelihood of this household making the observed choice is zero. Because, as we explain below, we use a maximum likelihood procedure to estimate our model, our model would not be identified without this assumption.

[24]This reflects the fact that there are no legal limits to the amount a household is allowed to spend on rent in The Netherlands. In practice, landlords do demand a certain minimum income. However, these demands are highly heterogeneous and not necessarily binding in practice. This is reflected by the fact that renters on average spend 42.2% of their disposable income on housing, where this is 26.6% for owner-occupiers (Statistics Netherlands, 2022).

counterfactuals. Endogenizing these factors by employing a dynamic model (Bayer et al., 2016) is possible, but this would add significant conceptual and computational complexity. As we have reasons to believe that the bias of using a static model is negligible for the our counterfactuals, we do not consider it worth the costs make the model dynamic.[25]

We close the demand model by making parametric assumptions on the distribution of the coefficients $(\alpha_{it}, \beta_{it})$ and idiosyncratic error terms $(\varepsilon^R_{ijt}, \varepsilon^O_{ijt})$. For the coefficients, we assume that they depend linearly on household characteristics as follows:

$$\begin{pmatrix} \alpha^R_{it} \\ \alpha^O_{it} \\ \beta_{it} \end{pmatrix} = \begin{pmatrix} \mu^{\alpha^R} \\ \mu^{\alpha^O} \\ \mu^{\beta} \end{pmatrix} + D'_{it} \begin{pmatrix} \gamma^{\alpha^R} \\ \gamma^{\alpha^O} \\ \gamma^{\beta} \end{pmatrix},$$

where D_{it} contains household characteristics and (μ, γ) are coefficients. For the error terms $(\varepsilon^R_{ijt}, \varepsilon^o_{ijt})$, we assume that they follow an Extreme Value Type I distribution, so that choice probabilities have the familiar logit form:

$$d^o_{ijt} \equiv P(i \text{ chooses house type } j,\ \text{tenure } o) = \begin{cases} \frac{\exp\{u^o_{ijt}\}}{1+\sum_k \exp\{u^R_{ikt}\} + \sum_{k\in\mathcal{C}_{it}} \exp\{u^O_{ikt}\}} & \text{if } k \in \mathcal{C}_{it} \text{ or } o = R, \\ 0 & \text{otherwise.} \end{cases}$$

Market-level demand follows from adding the choice probabilities of all households:

$$D^o_{jt} = \sum_i d^o_{ijt}$$

for all j and $o \in \{R, O\}$.

Supply and equilibrium. The available number of houses of type j is exogenous and given by S_{jt}. Hence, our model does not feature construction. However, the percentage of rentals is determined endogenously. In particular, we assume that for every house type there is a perfectly competitive market of buy-to-let investors. In equilibrium, the return on letting a house, p^R_{jt}/p^O_{jt}, must equal the hurdle rate of capital c. The hurdle rate is the equivalent to households' user cost of housing. Like for the user cost, it reflects the investor's cost of capital as well as additional benefits (like future appreciation) and costs (like maintenance). We assume this hurdle rate is heterogeneous, with distribution function $F_{jt}(c)$. In equilibrium, the fraction of rentals of housing type j, ϕ_{jt}, must then satisfy

$$F_{jt}\left(\frac{p^R_{jt}}{p^O_{jt}}\right) = \phi_{jt}. \tag{5}$$

To see the importance of the heterogeneity assumption, assume instead that all investors have the same opportunity hurdle rate $\bar{c}$. Equilibrium now requires that $p^R_{jt}/p^O_{jt} = \bar{c}$: the marginal return of investment in the rental sector equals the marginal cost. As a result, the rent-price ratio is exogenous and the supply of rental houses is perfectly elastic. As

[25] A second potential issue that the dynamic approach of Bayer et al. (2016) could solve is the correct computation of willingness to pay for attributes that change over time, such as neighbourhood crime. Since we focus on the willingness to pay for house characteristics that are constant over time, such as size and location, this is not an issue for our application.

Rouwendal et al. (2023) show, under this assumption, the equilibrium matching between households and house types is the same as in a market where consumers have unconstrained choice sets. Without the heterogeneity assumption, we would thus place a strong restriction on the supply of rental houses.

Heterogeneity of buy-to-let investors is plausible in the Dutch context as there are different types of landlords. In the market-rate sector, houses are let by by non-profit housing associations and private investors, although the former are more active in the rent-controlled sector. Private investors comprise both individuals and companies. In 2018, 47% of private landlords were individuals, with the rest being companies. Of the individual landlords, 80% only owned one rental property (Statistics Netherlands, 2019). There are hence large differences in profit motive, level of professionalism and ownership structure of market-rent landlords. It is likely that these differences translate into differences into the required rate of the return on housing. For example, institutional investors might have easier access to credit markets and hence a lower cost of capital.

To empirically pin down the distribution of investor heterogeneity, we assume that it takes the same shape for every house type j but that the distribution shifts up or down by a house type-specific factor η_{jt}. This factor reflects differences between housing types in maintenance costs, vacancy rates, growth expectations and risk premia.[26] This means that we write

$$F_{jt}(c) = F_t(c - \eta_{jt}), \tag{6}$$

where $F_t(\cdot)$ is the distribution function common across all house types in year t. This assumption implies that buy-to-let investors do not invest *across* different housing types. That is, if the return for house type j drops below an investor's required rate of return, this investors stops investing all together rather than switching its investment to a different house type. This is clearly a simplifying assumption, but we think it has some merit. Recall from the discussion above that a significant fraction of investors are private individuals, that own only a single property. These investors are relatively unsophisticated and hence less likely to invest in houses that are, for example, in a location they are unfamiliar with. Because these small-scale investors have worse access to capital markets than institutional investors, they are likely also the marginal investor. Hence, on the margin, it may be well be that investors drop out the market all together when there is a local shock to the return on rental properties.

Having derived the demand for housing and the supply of rental houses, we can now define equilibrium. In addition to the marginal rate of returns being equal to the investors' hurdle rate for every house type, the demand for houses must equal supply. Since there is separate demand for renting and owning a given housing type, there is a market clearing condition for both. The supply side of these conditions depends on behavior of buy-to-let investors through the no-arbitrage condition (5), with the number of rentals increasing at the expense of owner-occupied dwelling as buy-to-let investors become more active. This leads to the following definition of equilibrium.

Definition 1 (Equilibrium) *An equilibrium consists of a vector of house prices* $(p^O_{jt})_j$*, a vector*

[26]Campbell et al. (2009) and Bracke (2015) present evidence that rent-price ratios indeed vary with these factors.

of rents $(p_{jt}^R)_j$ *and a vector of rental probabilities* $(\phi_{jt})_j$ *such that for every* j*,*

1. $D_{jt}^R(p_t^R, p_t^O) = \phi_{jt}(p_{jt}^R, p_{jt}^O)S_{jt}$ *and* $D_{jt}^O(p_t^R, p_t^O) = (1 - \phi_{jt}(p_{jt}^R, p_{jt}^O))S_{jt}$ (market clearing*);*
2. $F_{jt}\left(\frac{p_{jt}^R}{p_{jt}^O}\right) = \phi_{jt}(p_{jt}^R, p_{jt}^O)$ (no arbitrage*).*

4.2 Identification and estimation

Estimation proceeds in two steps. First, we estimate the demand function from observed housing choices. Then, we use the estimated demand function to estimate investors' hurdle rate distribution.

Demand. We build on Bayer et al. (2004, 2007) to estimate household preferences for housing. They develop an efficient two-stage estimation procedure, in which the first stage estimates preference heterogeneity, while the second stage estimates the mean effects of house characteristics on utility. To do this, it is useful to decompose households' utility (4) as

$$u_{ijt}^o = \delta_{jt}^o + \lambda_{ijt}^o$$

for $o \in \{R, O\}$, where δ_{jt}^o is the utility common to all consumers and λ_{ijt}^o is idiosyncratic:

$$\begin{aligned} \delta_{ijt}^o &= X_{jt}'\mu^\beta - \mu^{\alpha^o} p_{jt}^o + \xi_{jt}^o, \\ \lambda_{ijt}^o &= X_{jt}'D_{it}'\gamma^\beta - D_{it}'\gamma^{\alpha^o} p_{jt}^o + \varepsilon_{ijt}^o. \end{aligned} \tag{7}$$

In the first estimation step, the heterogeneity parameters γ and the mean utilities δ are estimated by maximum likelihood. The log-likelihood is

$$\mathcal{L} = \sum_t \sum_i \sum_j \left(d_{ijt}^R \log p_{ijt}^R + d_{ijt}^O \log p_{ijt}^O \right),$$

where d_{ijt}^o is a dummy that indicates whether household i lives in house type j in year t, with tenure choice o. The parameter vector (δ, γ) is high-dimensional because there are two mean utilities, δ_{jt}^R and δ_{jt}^O, per house type j. Bayer et al. (2004) provide an efficient algorithm to maximize the log-likelihood. Their main insight is that, for all values of γ, it is possible to adjust the mean utilities δ such that aggregate demand exactly equals supply. Hence, it is possible to maximize the log-likelihood with respect to just γ, solving for the associated values of δ at every candidate value. Bayer et al. (2004) provide a contraction mapping to do so.[27] We refer the reader to their article for more details.

In the second estimation step, the mean utility parameters μ are estimated. To do so, note that (7) has the form of a linear regression. The left-hand side, δ_{ijt}^o, has been estimated in the first stage. The right-hand side contains only observable housing characteristics (X and p) and coefficients. ξ_{jt}^o is the econometric error term. However, estimating this regression using least squares would lead to inconsistent estimates of the price coefficients as prices are likely correlated with unobserved demand shocks ξ_{jt}^o.

[27] The contraction mapping provided in Bayer et al. (2004) is based on the assumption that the supply of every house equals one, which mirrors their setting but not ours. We hence modify the contraction mapping to allow for arbitrary supply, which is a straightforward extension.

It is therefore common in the literature to instrument for p, typically by using the characteristics of houses in nearby markets (Bayer et al., 2007). There are two problems with this approach in our context. The first is that the housing characteristics we have in our data set, such as square footage and year of construction, do not have a lot geographical variation. As a result, there is a risk such instruments are weak in our setting. Moreover, there is a general concern that product characteristics sometimes yield inconsistent estimates (Armstrong, 2016). Therefore, we depart from the literature by not using instruments for price.

Instead, we exploit the fact that we have *two* prices per house type, the purchase price and the rent, to identify the price coefficients. The idea is that demand shocks affect owner-occupied and rented houses of the same type similarly. By comparing the demand response to changes in the rent-price ratio between owner-occupied and rental houses of the same type, we can then control for these shocks. Formally, we make the following identifying assumption:

$$\xi^o_{jt} = \bar{\xi}_{jt} + \tilde{\xi}^o_{jt}, \ \mathbb{E}[\tilde{\xi}^o_{jt} | p, X] = 0, \tag{8}$$

for all j, t and $o \in \{R, O\}$.

Substituting this decomposition of the error term into (7) gives

$$\delta^o_{jt} = \bar{\xi}_{jt} - \mu^{\alpha^o} p^o_{jt} + \tilde{\xi}^o_{jt} \tag{9}$$

for all j, t and $o \in \{R, O\}$. Hence, we can identify the mean price coefficients with a fixed effects regression. The term relating to product characteristics X_{jt} has dropped since we cannot identify it separately from the housing type-year fixed effect $\bar{\xi}_{jt}$. Therefore, we cannot identify the effect of house characteristics other than price on mean utility. Since we do not vary these characteristics in our counterfactuals, this is no limitation in our application.[28]

The formulation (9) further elucidates our identification strategy. Since we control for house type-year fixed effects, we identify μ^{α} through changes in the rent-price ratio of a given house type over time. For this strategy to be valid requires that changes in the rent-price ratio are driven by supply and not by demand. We now provide some suggestive evidence for this assumption. First, buy-to-let investors are likely more responsive to shocks than owner-occupiers. This is because they face lower transaction costs as they do not have to move. As a result, shocks that impact the rent-price ratio operate through the supply side more than through the demand side. Second, we exploit that, while positive demand and negative supply shocks can both lead to price increases, they have opposite effects on quantities: a positive demand shock increases the equilibrium quantity, while a negative supply shock decreases it. Therefore, if changes to the rent-price ratio are primarily caused by supply shocks, we expect a negative correlation between the price/rent ratio and the fraction of rentals. We test this implication in Table 3, where we regress the fraction of rentals on the price/rent ratio and house type-fixed effects. We indeed find a negative correlation, consistent with supply shocks being more important than demand shocks.

[28]If the effects of house characteristics, X, are of special interest, the fixed effects $\bar{\xi}_{jt}$ could be decomposed further, as in Nevo (2001).

Table 3: Importance of supply vs. demand shocks

Dependent Variable:	Fraction rentals
Variables	
Rent-price ratio	-0.2520***
	(0.0394)
Fixed-effects	
House type	Yes
Fit statistics	
Observations	29,325
R^2	0.96443
Within R^2	0.01175

Note: The table shows the coefficients of a regression of the fraction of rentals on the price/rent ratio and house type fixed effects, 2015–2019. When supply shocks are more important in determining shocks to the rent/price ratio, we expect a negative coefficient. When demand shocks are more important, the coefficient should be positive. See the main text for further explanation. Clustered (House type) standard-errors in parentheses. Signif. Codes: ***: 0.01, **: 0.05, *: 0.1

The third piece of evidence for our identification strategy are our results. Based on the literature, we have a reasonable idea what the user cost of housing should be (Poterba, 1984). As we show below, we indeed find reasonable values for the user cost of housing.

The identifying assumption (8) implies that there are no structural quality differences between owner-occupied and rental houses. In other words, demand differences between owner-occupied and rental houses of the same type are attributed completely to differences in price sensitivity. This is contrary to some theoretical results (e.g. Henderson and Ioannides, 1983) that claim that rental houses are of lower quantity due to misaligned maintenance incentives between the landlord and the occupant. To correct for this, we tried to additionally control for Rental $\times$ House type fixed effects in regression (9). However, as the regression results in Supplementary Material E show, we do not have sufficient variation in our data to disentangle price sensitivity from quality differences: we obtain very large standard errors on the price coefficients. Rather than continue with an imprecisely estimated demand curve, we use the main specification (9). As a result, we underestimate the user cost of housing to the extent that rentals are of lower quality than observationally similar owner-occupied houses. As we discuss below, we think our estimated user cost of housing is in a reasonable range. We hence believe that this bias to the user cost of housing is small. Additionally, we show in Supplementary Material E that our main results are qualitatively similar using a specification *with* Rental $\times$ House type fixed effects. Our results are thus do not depend on our exactly disentangling tenure preferences from quality differences.

For the identification of the non-price parameters, we require variation in and/or between households' choice sets. Intuitively, comparing the choices of households with different options gives information on the value they attach to these options. Typically, this variation is obtained by comparing different (geographical) markets. We take a different approach. Our setting provides significant exogenous variation through DSTI limits. DSTI limits differ between households, creating cross-sectional variation in choice sets. But, as we already noted in Section 3, there is also exogenous temporal variation in DSTI limits. Hence, DSTI limits create exogenous differences between households' choice sets, which identify the non-price parameters of the model.

We use the following household and house characteristics to model heterogeneity. For households, we use the average age of adults in the household, the number of persons in the household, log household income and household wealth. We distinguish between houses based on their square footage, the urbanity of their location, and a dummy indicating whether it is an apartment of a singly-family property. As our house types include multiple houses, we take the average over these characteristics, e.g. for every house type we have a variable that includes the fraction of apartments in this type. However, as we show in Supplementary Material C, our house types are relatively homogeneous in these characteristics. We similarly take the average rent and purchase price per house type. Since the price sensitivity of owner-occupied properties depends on the interest rate, we interact the purchase price with the average mortgage interest rate in both the first and second stages of our estimation.

Supply. On the supply side, we need to estimate investors' hurdle rate distributions $F_{jt}(\cdot)$. To obtain an estimable equation, combine (5) and (6) and solve for the rent-price ratio write

$$\frac{p_{jt}^R}{p_{jt}^O} = F_t^{-1}(\phi_{jt}) + \eta_{jt}. \tag{10}$$

The rent-price ratio p_{jt}^R/p_{jt}^O and the fraction of rentals ϕ_{jt} are directly observable in the data. Hence, we can estimate the inverse cdf of the hurdle rate distribution, $F_t^{-1}(c)$, with a non-parametric regression.

Here, we face an identification challenge because the fraction of rentals ϕ_{jt} is likely correlated with the error term. For example, if expected future rent appreciation increases, investors are willing to pay more for the house for any given level of demand. The rent-price ratio then decreases while the fraction of rentals increases. To control for this type of endogeneity, we need an instrument for the fraction of rentals that is uncorrelated with supply shocks.

We derive such an instrument from the demand side of our model. To be a valid instrument, the instrument needs to be correlated with the fraction of rentals but be uncorrelated with supply shocks to the rent-price ratio. We therefore build an instrument that measures the relationship between households' preference to rent and horizontal attributes of a house type. To do this, we compute demand from our estimated model in an environment with two changes: i) all prices are set to zero, ii) the mean utility of all houses (δ) is set to zero. We then compute the average user cost of housing per house type as a measure of the preference to rent. By setting the price to zero, we remove any influence of supply shocks that operate through the price. By setting the mean utility to zero as well, we also remove the influence of any demand shocks that are correlated with supply shocks.[29] What is left is demand that is only driven by household heterogeneity. Intuitively,

[29]For example, if the expected price appreciation of a house type increases, the demand for rentals shifts to the left, while the supply of rentals shifts to the right.

if a house type contains a lot of apartments, it will in our experiment be demanded by households that have a relatively large preference for apartments. If these households also have, say, a large preference to rent, this is an exogenous shock to the fraction of rentals. We stress that our instrument measures demand for a hypothetical, exogenous supply of houses and that we do not impose market clearing. It is hence uncorrelated with the supply side of the market by construction.

The validity of the instrument requires that the estimated differences between households' willingness to pay for housing attributes are uncorrelated with supply shocks to the rent-price ratio. A sufficient condition is that our model of preference heterogeneity for housing characteristics is valid. In particular, a consistent estimate of the correlation between tenure preferences and preferences for housing characteristics is sufficient. The estimated preference heterogeneity, and hence our instrument, depends on i) the selection of house and household characteristics included in the model, and; ii) the functional form of the demand function. This means our instrument is somewhat dependent on functional form. We have computed some variants of our demand model, where we have included higher-order terms of household characteristics. The coefficients on these terms were typically small and insignificant. While we cannot prove that our functional form is correct, this gives us some confidence that our model does not suffer from a large degree of misspecification.

Figure 4 shows our instrument aggregated to a municipality level. It shows that the average user cost is smaller in larger municipalities. This means that households that have a larger preference for the type of houses that can be found in larger cities, also have a relatively large preference for renting. As a result, the demand for rentals is larger in these municipalities and we should expect more rentals there for any level of supply. This suggests that our instrument is valid. In Table D3 in the Supplementary Material, we also compute the first stage of a linear version of (10). Here, we show that our instrument is also highly relevant.

We estimate our supply-side model (10) using the non-parametric instrumental variables (NPIV) estimator by Chetverikov and Wilhelm (2017). This estimator allows us to impose a monotonicity constraint on the inverse cdf $F_t^{-}(c)$. As shown by Chetverikov and Wilhelm (2017), this can greatly improve the performance of the estimator. Since the inverse cdf must be increasing, we have a natural monotonicity constraint to impose. The NPIV estimator uses B-splines to estimate the non-linear relationship in both the first and second stages. We pick the degree and number of knots of these splines using cross-validation.[30]

4.3 Computing counterfactuals

One of the advantages of our structural approach is that we can compute the counterfactual equilibria. As policy changes likely have large equilibrium effects through house prices and rents, reduced form approaches as in Section 3 are limited. A structural model also allows us to compute the effects of policy changes on consumer and investor surplus.

To compute an equilibrium, we numerically solve for prices and the fraction of rentals

[30]We are grateful to Chetverikov et al. (2018) for making a Stata package available for the NPIV estimator.

Figure 4: Supply-side instrument: relation between horizontal preferences for houses and preference for renting

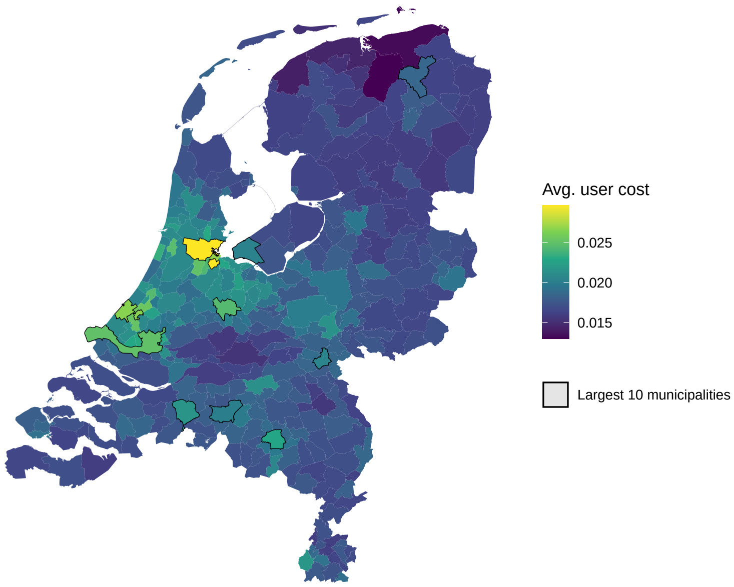

Note: The graph shows the average user cost of housing per municipality computed from a demand function where houses are only horizontally differentiated, for 2019. To compute demand when houses are only horizontally differentiated, we compute demand when the prices and vertical components of utility are both zero for all houses. The user cost measures the preference for renting. They are computed per house type and then translated to a municipality level for this graph. The average user cost is computed as the average of the ratio of price coefficients, $\alpha^O_{it}/\alpha^R_{it}$, weighed by the probability that a household lives in a certain house type.

per house type. That is, demand must equal supply for every house type and tenure choice, and the marginal buy-to-let investor must make zero economic profits. When doing so, we make one simplifying assumption on DSTI limits. In reality, the maximum a household may pay for a house depends on two factors: the maximum mortgage and its wealth. The former depends on DSTI limits and income. Therefore, it is exogenous. This is not true for household wealth. As it increases with the equity a household has in its house, it depends on equilibrium prices and is thus endogenous. These kind of feedback loops are a technical challenge as increases in house prices can both expand and shrink choice sets. Therefore, we ignore these feedback loops by holding household financial wealth constant in our counterfactuals. As a result of this restriction, we will tend to underestimate the magnitude of changes to house prices. However, this effect is likely to be small as home owners in our data set tend not be constrained by DSTI limits: 97% of home owners live in a house with a market price at least 10% below the maximum they could pay.

Having computed an equilibrium, we can determine changes to consumer and investor surplus. Here, we have to account for the fact that households and investors not only buy, but also sell houses. This differs from product markets where discrete choice type models are also commonly used, where there is a clear separation between consumers and producers. As we employ a static equilibrium concept, we take the following approach. In essence, we assume that all houses are on the market and are reallocated. Every household that chooses owner-occupation incurs the market price, while every household that owns a house in our data sells it and gains the equity that it has in the house. In other words, a household that doesn't move buys the house from itself. Total consumer surplus of household i can then be written as

$$CS_{it} = \mathbb{E}[WTP_{it} - P_{it}] + E_{it},$$

where $\mathbb{E}[WTP_{it}]$ is the expected willingness to pay of household i in year t, $\mathbb{E}[P_{it}]$ is its expected expense on housing services and E_{it} is the surplus it derives from the equity it has in its house. Renters' equity is zero. The expectation term is the money-metric-utility the household derives as a consumer in the housing market, while the last term is the utility it obtains as a potential seller.

The demand side utility follows from our estimated demand model. As is standard for discrete choice models, we can make expected utility money-metric by dividing it by the coefficient on price (Small and Rosen, 1981).[31] This works since the price coefficient measures the marginal utility of income. In our case, this is not strictly true. As (2) shows, the ratio on the coefficient rent and on utility is

$$\frac{\rho + \kappa}{\rho + \kappa + \kappa\psi - \psi}\alpha.$$

This is only equal to the marginal utility of income α, as assumed, when the real growth of rents ψ equals zero. This is because households do not only react to current rents in their housing choice, but also to expected future rents. When $\psi > 0$, the coefficient on rents overestimates the marginal utility of income. All our welfare calculations would then be biased away from zero. Real rent growth of existing contracts is limited in the Netherlands, as rents are typically indexed to inflation plus a small increment: during our sample real rents of existing contracts increased between 0 and 1.5% per year.[32] We hence believe bias of using the coefficient on rents to compute the marginal utility of income to be small.

Here, we face also the issue that our model has *two* price coefficients: one for rentals and one for owner-occupied dwellings. We therefore have to convert between a continuous payment stream, rent, and a one-time payment, a house's price. Here, we use the insight that the user cost of housing measures the marginal rate of substitution between housing and other consumption (Dougherty and Van Order, 1982). We can therefore convert the utility from owner-occupation to a yearly figure by dividing it buy the user cost of housing, ω_{it}.[33] We take the same approach when considering households' role as seller in the

[31]This is technically not correct when the coefficient on price is a function of income, as in our case. However, the mismeaurement induced by this misspecification is small when changes to consumer surplus are small relative to income (Train, 2009, p. 97). We have verified for a small subset of our sample that our computation of consumer surplus is similar to that that is obtained when we numerically solve for the equivalent variation, as in McFadden (2012).

[32]Nominal rents increased between 2% and 3% per year (own computations), while inflation was between 0.3% (in 2016) and 2.6% (in 2019).

[33]As we describe below, some households have a negative user cost of housing. In this case the relationship between the user cost and the marginal rate of substitution between housing and other consumption breaks down, as a this rate cannot be negative. We thus take the user of cost of housing as zero for these households

housing market: we compute E_i by discounting household's equity with their user cost.

Our derivation of investor surplus follows a similar logic. Investors receive a yearly stream of rents. In addition, they incur capital costs on and derive profits from price increases of properties they own. Hence, we can write the yearly surplus of investor i in year t as

$$IS_{it} = R_{it} - P_{it} + E_{it},$$

where R_{it} is the yearly rent it receives and P_{it} and E_{it} have the same meaning as for households. As for households, we have to discount one-time payments to purchase houses and changes into equity to a yearly stream to make them comparable with rental income. Analogous to the household case, the right discount factor for investors is their hurdle rate, c_{it}.

These decompositions of consumer and investor surplus help elucidate the sources of aggregate welfare gains and losses in our model. The first is measured by changes in total willingness to pay, $\sum_i \mathbb{E}[WTP_i]$. This measures the utility generated in the market, for example due to changes in choice sets. The second source of aggregate welfare effects follows from redistribution of capital costs. As the discount factor used to compute P_i and E_i is heterogeneous, it matters who buys and owns houses. For example, a given house price leads to a larger decrease in consumer surplus when it is incurred by a household with a larger user cost. This reflects that such a household has a greater marginal rate of substitution between housing and other consumption. Hence, price changes need not be welfare neutral in our model. This result contrasts with previous finding in the literature that, since every buyer is also a seller, changes in house prices have no effect on aggregate welfare (Bajari et al., 2005).[34]

We hence measure allocative efficiency in the housing market instead of all possible effects of DSTI limits. In particular, we do not consider any externalities of DSTI limits. Such externalities can be direct, e.g. DSTI limits can have positive effects on financial and macroecnoomic stability (Claessens, 2015). There are are also indirect externalities through the effects of DSTI limits on housing market equilibrium. For example, if DSTI limits increase the fraction of rentals, this can have an effect on the labour market as renters are more likely to take a job in a different city (Head and Lloyd-Ellis, 2012). As is the case for any welfare measure, total surplus also cannot reflect paternalistic motives to protect households from high debts.

5 Debt limits and buy-to-let investors

We now use our model to understand the relationship between debt limits and buy-to-let investors. We start by explaining the mechanisms behind tenure choice in our counterfactuals. Then we compute to what extent changes in DSTI limits can account for the observed changes in rentals in the Netherlands. We finish by discussing the effects of changes to DSTI limits, as well as other policies that deter buy-to-let investing such as rent taxes.

when computing their surplus.

[34]We note that this result would be obtained in our model if all user and capital costs were equal.

Figure 5: Distribution of user cost of housing

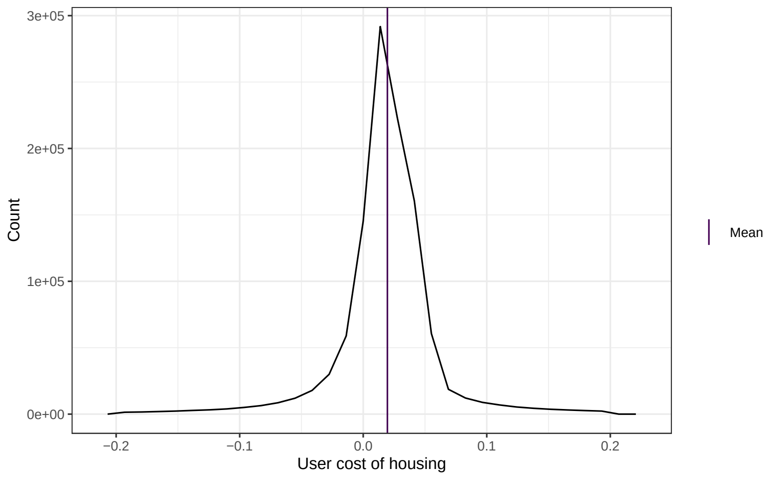

Note: The figure shows the estimated distribution of the user cost of housing for 2019. The user cost of housing follows from our estimated demand model, see the main text for more details. Values below -.2 and above .2 are not shown for clarity.

5.1 The user cost of housing

As the user cost of housing is central to households' tenure choice, we start by discussing our estimates of its distribution. We leave discussion of other model primitives to Supplementary Material E. Figure 5 shows our estimated distribution for 2019. The mean user cost of housing is 1.95%. The average mortgage interest rate after tax deductions was 1.27% in 2019.[35] The user cost of housing should be somewhat higher than the mortgage interest rate, as it additionally measures expected maintenance, price risk, etcetera (Poterba, 1984). We find a gap of 60 basis points between the average user cost of housing and the average mortgage interest rate. We believe this be a reasonable difference, which shows support for our identification strategy.

Surprisingly, 20% of households have a negative user cost of housing. This means that they prefer to own their house even if they could live in it rent-free. We believe this reflects the treatment of housing equity in the Dutch tax system. Housing equity of primary residences is not taxed, while other forms of financial wealth are typically taxed at a rate of 1.2%. This means that for households with significant wealth, owning a house can be more profitable than living in it rent-free. Consider, for example, a household that buys a €600,000 house in cash. By doing so, it saves €7,200 in wealth taxes per year. If the risk-adjusted return on housing is as good as on other investments, the household would hence have to be paid at least €7,200 per year not to buy the house. In other words, for wealthy households with owner-occupation is almost always preferable to renting, which is reflected by a negative user cost of housing. Consistent with this story, we indeed find that households with negative user costs are significantly wealthier than households with positive user costs.[36]

[35] The average interest rate on 10-year mortgages was 1.99%, while most households deduct mortgage interest at a 36.1% income tax rate. This gives the net interest rate of 1.27%

[36] Households with a negative user cost have a median wealth of over €490,000. For households with a positive user cost, this is just under €130,000.

Figure 6: Illustrative equilibrium without housing heterogeneity and capital constraints

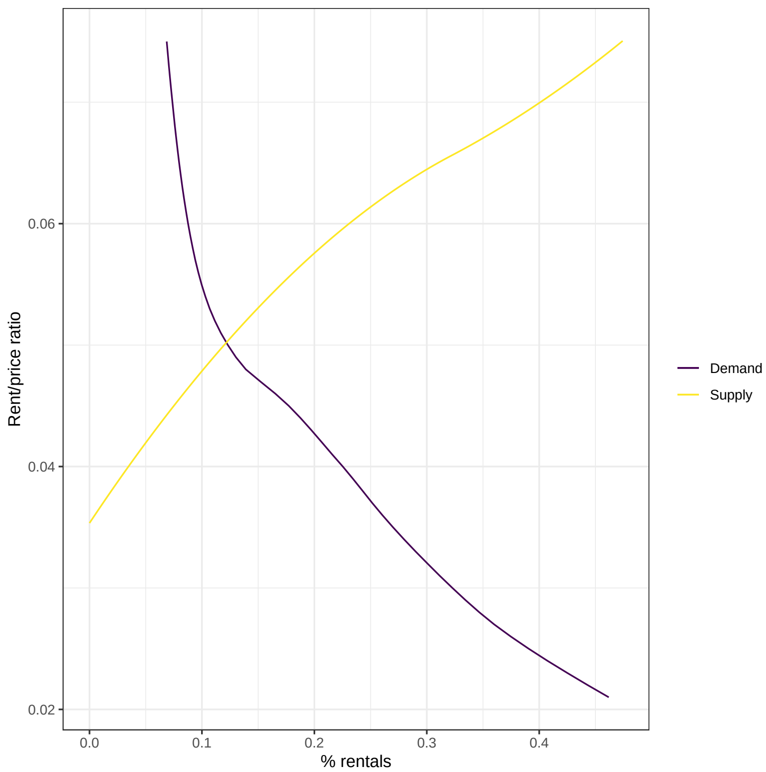

Note: The Figure shows an illustrative equilibrium for an economy when houses are homogeneous and households do not face DSTI limits. The supply curve is the estimated cdf of investors' hurdle rate distribution ($F_t(\cdot)$ in our model). The demand curve is the empirical survival function of households' user cost of housing (the ratio of price coefficients in our model). Both graphs are computed for 2019.

5.2 The effects of DSTI limits on tenure choice

If households are unconstrained, their demand for rentals depends on their user cost of housing. Whether a house is a rental in equilibrium, however also depends on investors' hurdle rate. If the user cost of housing is greater than investors' hurdle rate, the house will become a rental. This is the case since the investor is willing to pay more for the same house. Hence, if it were owner-occupied, a mutually beneficial transaction could take place where the household sells the house to an investor and rents it. Similarly, the house will be owner-occupied whenever the household's user cost of housing is below the investor's hurdle rate.

To gain some intuition into this mechanism, we start by exploring equilibrium in an economy where i) houses are homogeneous, and ii) households are unconstrained. In such an economy, the demand for rentals is equal to the number of households that has an user cost *above* the rent/price ratio. In other words, the demand for rentals is the survival function of the user cost of housing.[37] The supply of rentals follows from the number of investors that have a hurdle rate *below* the rent/price ratio. Hence, the supply of rentals equals the distribution function of the hurdle rate distribution. We plot these distributions in Figure 6, flipping the axes to obtain a plot of supply and demand. Equilibrium occurs whenever supply and demand interact, so that the rent/price ratio equals the hurdle rate of the marginal investor.

[37]The survival function equals 1 minus the distribution function.

Figure 7: Partial equilibrium for 3 house types, with and without DSTI limits

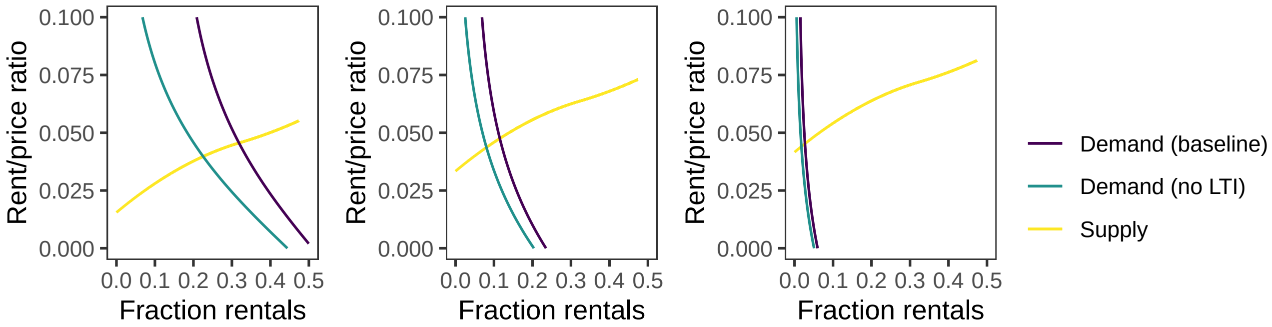

Note: The Figure shows the partial equilibrium for three representative housing types: one where the equilibrium number of rentals is high (left), one where it is comparable to the national average (middle) and one where it is lower (right). The demand curves are computed by computing for every household the total probability that it lives in a house of that type. Holding these probabilities constant, we vary its rent (holding the purchase price constant) to compute the probability that households live in a rental of this type. To obtain the demand curve when households would face no DSTI limits, we follow the same exercise but allow for an unrestricted choice *within* the housing type. Supply curves are estimated using non-parametric instrumental variables and are adjusted for a housing-type fixed effect compared to Figure D1.

In our model, however, the number of rentals differs between house types. There are three causes for this. First, preferences for housing characteristics are correlated with the user cost of housing. As a result, households with different tenure preferences self-select into different house types. Second, households are constrained by DSTI limits. The demand for rentals is thus not only determined by households' preferences, but also by the extent to which DSTI limits bind. Third, the hurdle rate distribution can shift to the left or to the right per house type, so that the supply of rentals shifts inwards or outwards.

To see the effects of these factors, we show the partial equilibrium for three house types (Figure 7). Here, we compute the demand curve by holding the total probability that a household lives in the housing type constant, then varying the rent. We compute demand this way with and without DSTI limits. We can obtain a couple of insights from these supply and demand graphs. First, the constrained efficient number of rentals, which is the intersection of the supply curve and the unconstrained demand curve, differs between housing types. This reflects a combination of households' self-selection and differences in the hurdle rate distribution between housing types. Second, DSTI norms distort households' tenure choice by shifting the demand curve for rentals to the right. However, the size of the distortion differs between housing types, as the extent to which households are constrained is correlated with preferences for housing attributes.

DSTI limits hence distort households' tenure choice, but the question remains to what extent the tightening of DSTI limits explains the observed increase in rentals in the Dutch housing market. To quantify this effect, we take our estimated model from 2019 and recompute an equilibrium under 2013's DSTI limits. We take 2013 as our comparison year, as most of the changes in DSTI limits happened after this time (Figure 1a). Figure 8 shows the results from this exercise. Relaxing DSTI limits to 2013 levels would decrease the fraction of rentals by 8.65%. Since the fraction of rentals has in fact increased by 70% between 2013 and 2019 (Figure 1a), changes in DSTI limits in this period can explain 21%

Figure 8: Effects of reversing 2013–2019 tightening of DSTI limits

(a) Effects on prices and rentals

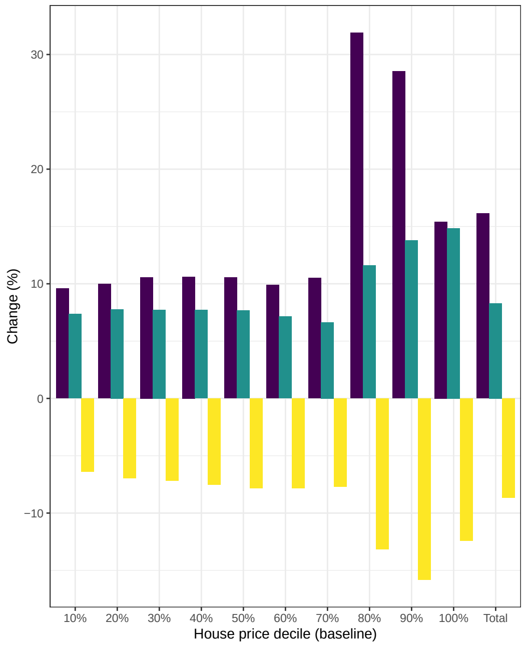

(b) Geographical variation

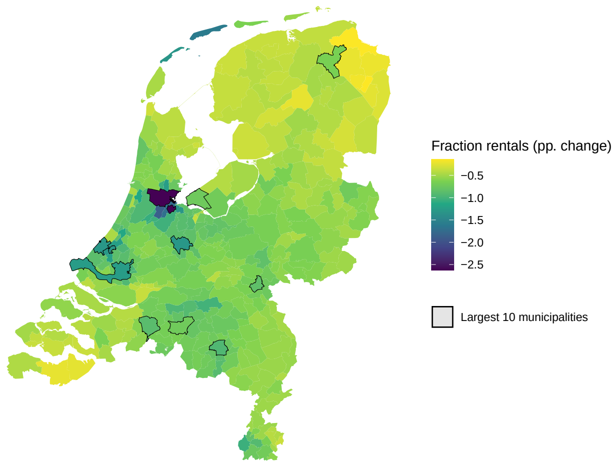

Note: The figures show a counterfactual where DSTI limits have been relaxed to 2013 levels. Counterfactuals have computed by taking the estimated 2019 equilibrium and imposing 2013's DSTI limits. The left graph shows the effects on equilibrium prices by government-assessed house price. The right graph shows the effect on the fraction of rentals by municipality.

of the increase in rentals.[38] Therefore, DSTI limits cause a significant increase in rentals, but are far from sufficient to explain all changes.

These estimates imply an elasticity of the equilibrium fraction of rentals with respect to debt limits of -0.32. This estimate is of a comparable order of magnitude as the average elasticity we computed in our reduced-form results (-0.22). In our view, this shows that our structural model is able to capture the main mechanisms through which debt limits influence tenure choice.

As can seen in the right panel of Figure 8, DSTI limits primarily caused an increase in rentals in larger municipalities. This makes sense as larger municipalities attract relatively young households, that are more likely to be constrained by DSTI limits as they do not have equity in their current home. Driven by the more liberal DSTI limits, house prices would increase by 15% on average. This would be most pronounced in the upper third of the housing price distribution, with increases close to 30% as the demand for owner-occupied housing increases. Prices on the rental market are affected in a more indirect way. While increased demand for owner-occupied housing would decrease demand for rentals at first, higher house prices mitigate the effect and steer people back to the rental market. As buy-to-let investors sell part of their property stock at higher house prices, supply of rental properties decreases. In our counterfactual these countervailing forces dominate the initial demand shift to owner-occupation, resulting in an increase of rents by 8%.

[38] To reverse the 70% increase, rentals would need to decrease by $1 - 1/1.7 = 41.2\%$. Since, in our counterfactual, rentals decrease by 8.65%, DSTI limits can explain $8.65/41.2 = 21\%$ of the observed increase.

5.3 Policies to reduce the distortion of DSTI limits

A natural follow-up question to the observation that DSTI limits distort households' tenure choice, is whether policy interventions to reduce the number of rentals can be justified. When answering this question, it should be remembered that DSTI limits are not the only distortion in the housing market. For example, owning a home is subsidized in the Netherlands, as in many other countries. Hence, even though DSTI limits distort the market toward rentals, it is not obvious that the number of rentals is too high when we take all distortions into account. Since our model takes all distortions other than DSTI limits as given, we can only speak to constrained efficiency. Nevertheless, we think this a worthwhile exercise. First, other distortions may serve a particular policy goal, e.g. correcting for some externality. In this case, constrained efficiency is the right welfare measure. Second, even absent such a justification, we observe that primarily lower-income households rent (Figure 1b). Governments may view it as undesirable when these households are excluded from owning a home for redistributive reasons. We also net out the effect of subsidies in our welfare measures, as we explain below.

We therefore compute the effects of various policy measures to combat the distortion on tenure choice introduced by DSTI limits. The most obvious measure is to relax DSTI limits. However, governments may be reluctant to do so because they want ensure the stability of the financial system. They may also want to protect households against excessive borrowing out of paternalistic motives. Our model cannot evaluate these reasons. We thus restrict our analysis to the direct effects of DSTI limits on housing market equilibrium.

As Figure 7 shows, the problem is that DSTI limits distort households' tenure choice towards renting. Their actual demand is therefore to the right of their welfare-relevant demand curve, which is obtained by removing all DSTI limits. The goal is hence to shift the equilibrium to the left. The first possible measure is an introduction of a tax on rental income, which would shift supply to the left.[39] Tax increases on investment properties have been introduced in the Netherlands and the UK. For example, the UK has an additional 3% stamp duty for investment properties, with a further 2% duty if the investor is foreign.[40] The Netherlands had a similar increase in stamp duty for investors from 2% to 10.4% and has also increased property taxation on investment properties.[41] (Both measures were taken after the period we consider.) A second measure is to cap the quantity of rentals. This reflects a measure taken in the Netherlands, where municipalities can ban investors from purchasing houses.[42] This is in practice only done for house types and areas with a large number of rentals, so that it forms an implicit cap on the quantity of rentals. Some countries only ban certain types on buy-to-let investors, which amounts to a partial cap. For example, Canada bans foreign buy-to-let investors.[43] Finally, a cap on rents also reduces to the number of rentals. Rent control is common to many countries, but unfortunately we

[39] Our supply model is estimated given taxes on rent that are already present in the Dutch context. Technically, we are thus computing an *increase* in the taxation of rental income.

[40] https://www.gov.uk/stamp-duty-land-tax/residential-property-rates. Accessed May 1, 2023.

[41] https://www.rijksoverheid.nl/onderwerpen/belastingplan/vermogen-en-wonen/overdrachtsbelasting and https://www.rijksoverheid.nl/onderwerpen/belastingplan/vermogen-en-wonen/leegwaarderatio. Accessed May 1, 2023.

[42] https://www.volkshuisvestingnederland.nl/onderwerpen/opkoopbescherming. Accessed May 1, 2023.

[43] https://www.bbc.com/news/world-us-canada-64082923. Accessed May 1, 2023.

are not able to compute its effects in our model.[44]

All these measures restore the constrained efficient equilibrium when they are house type-specific. For example, in Figure 7, a larger rent tax is required for the left-most housing type than for the right-most. But a house-type specific tax is not feasible in practice. We thus consider the other extreme of a uniform tax. In practice, governments may choose different tax rates for broad segments of the housing market, for example cheap and expensive houses. Our results provide some guidance whether uniform rates are at all able to achieve welfare gains. Like for the rent tax, the optimal cap on the fraction of rentals is house type-specific, but we consider a uniform cap. To be precise, we impose the a cap per house type. The idea is that governments have some (implicit) policy goal to constrain the number of rentals and start to impose measures when this number is reached in some locality.

As one of policy makers' concerns around the effects of DSTI limits is the exclusion of households from purchasing a home, we start with the effects of our policies on the number of constrained renting households. As most of these households are in the bottom 30% of incomes of our sample (Figure 1b), we focus on this subsample. All policies except DSTI decreases reduce the number of constrained renters (Figure 9), but not as much as one perhaps would expect based on our reduced-form results. Among changes to DSTI limits, the largest effect we obtain is from a decrease in DSTI limits of 10%. We find that the fraction of constrained renters then decreases with 1.12 percentage points (11.4%), while a naive extrapolation of our elasticity would yield a decrease of 100%. The main reason for the differences in effect sizes is that our counterfactuals are able to account for the equilibrium effects of debt limits on prices, while our reduced-form results measure the effect of a ceteris paribus change to debt limits to a small portion of the population. This equilibrium effect partially undoes the policy's direct effect on constrained households' choice sets.

We now compute the policies' effects on total surplus. When computing total surplus, we have to confront the fact that our estimates of households' willingness to pay include the effects of fiscal policies. Since it is subsidized, a higher willingness to pay for owner-occupancy can reflect an intrinsic preference or the effect of subsidies. We net out these transfers as follows. First, we assume that if a household moves from renting to owning it takes out a mortgage of 100%. We can then compute the increased subsidy from the mortgage interest deduction.[45] The net fiscal subsidy to owner-occupancy is about 20% higher than the mortgage interest deduction (Ewijk and Lejour, 2017). We hence multiply the mortgage interest deduction with a factor 1.2 to obtain the net fiscal subsidy. Since we conservatively assume that all additional expenses on home purchases are fully financed by mortgage, we are likely to somewhat overestimate the size of fiscal subsidies. Our welfare effects should hence be viewed as a lower bound.

Figure 10 shows the effects of our policy measures on total surplus. We decompose

[44]The reason is that under rent control, the equilibrium features rationing. Since households are heterogeneous in our model, a crucial question is then which households are rationed. But we do not know the rationing rule that would occur under rent control. Since the rationing rule can be of first-order importance to the welfare effect of rent control (Glaeser and Luttmer, 2003), we run the risk of in essence assuming our results by assuming a certain rationing rule.

[45]We assume a uniform interest rate of 1.99%, which was the average rate on 10-year mortgages in 2019.

Figure 9: Effects of policies on constrained households

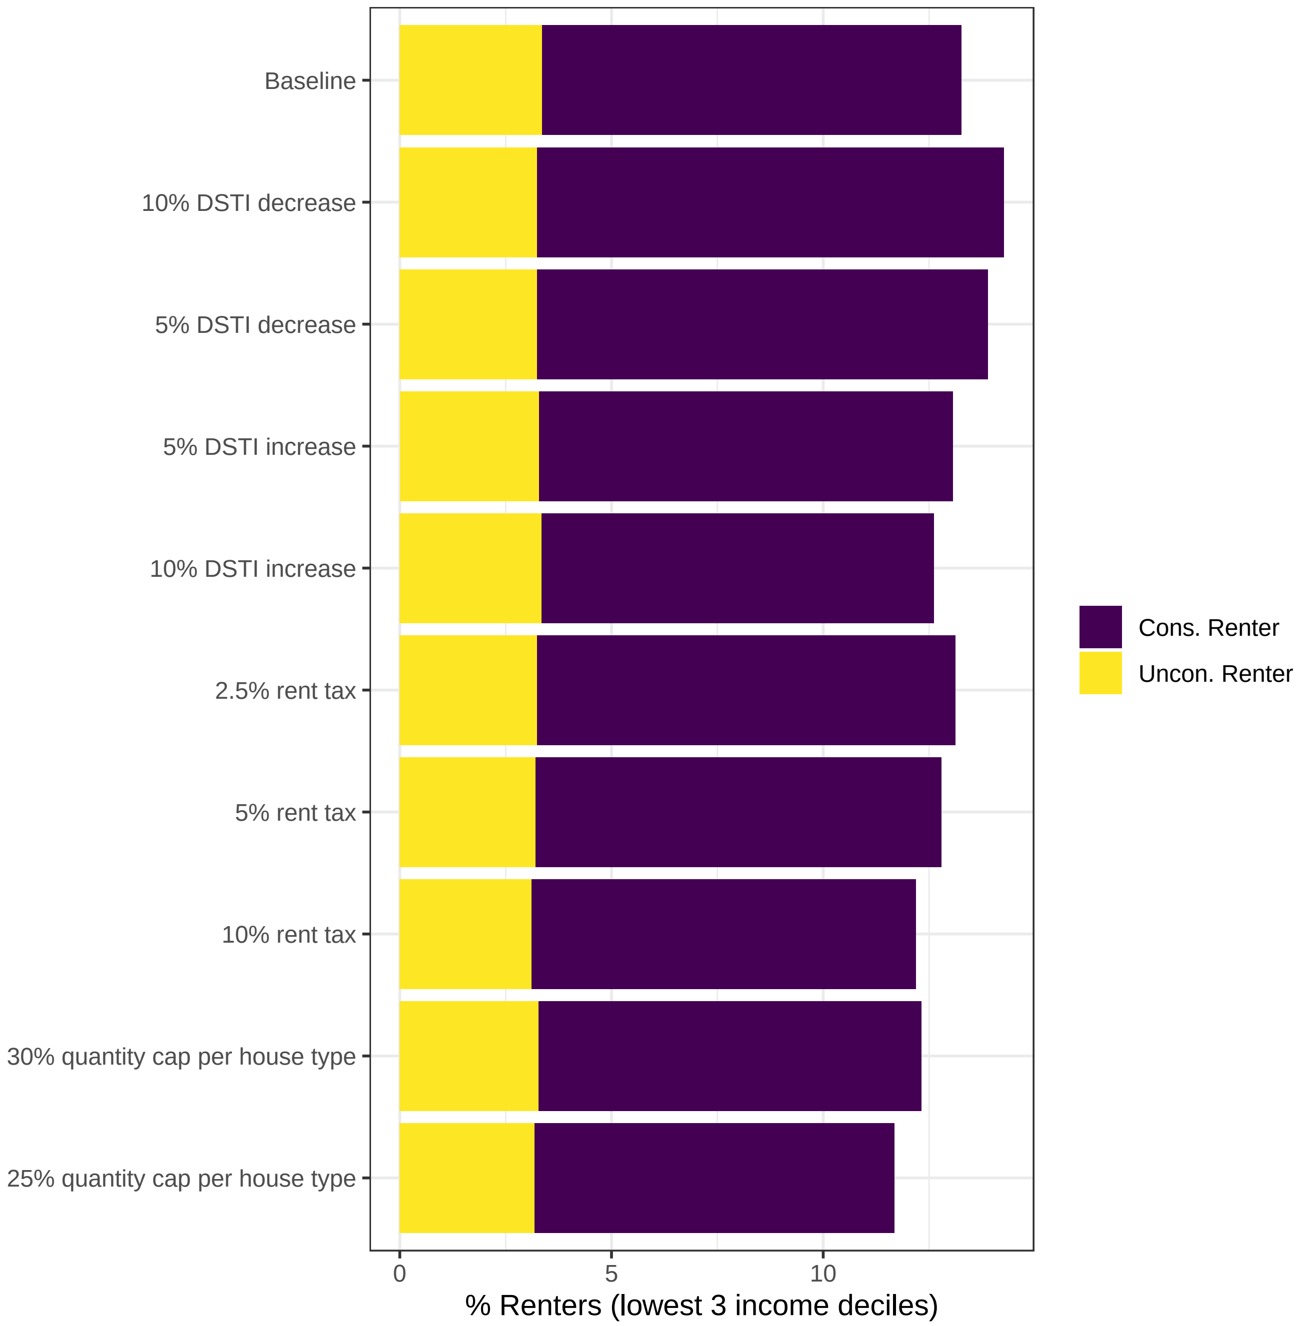

Note: The figure shows the fraction of renters across households with household incomes in the bottom 30% of our sample for different counterfactuals. Constrained renters are renters that cannot purchase a house of the same type that they are predicted to rent due to DSTI limits.

Figure 10: Effects of policies on total surplus

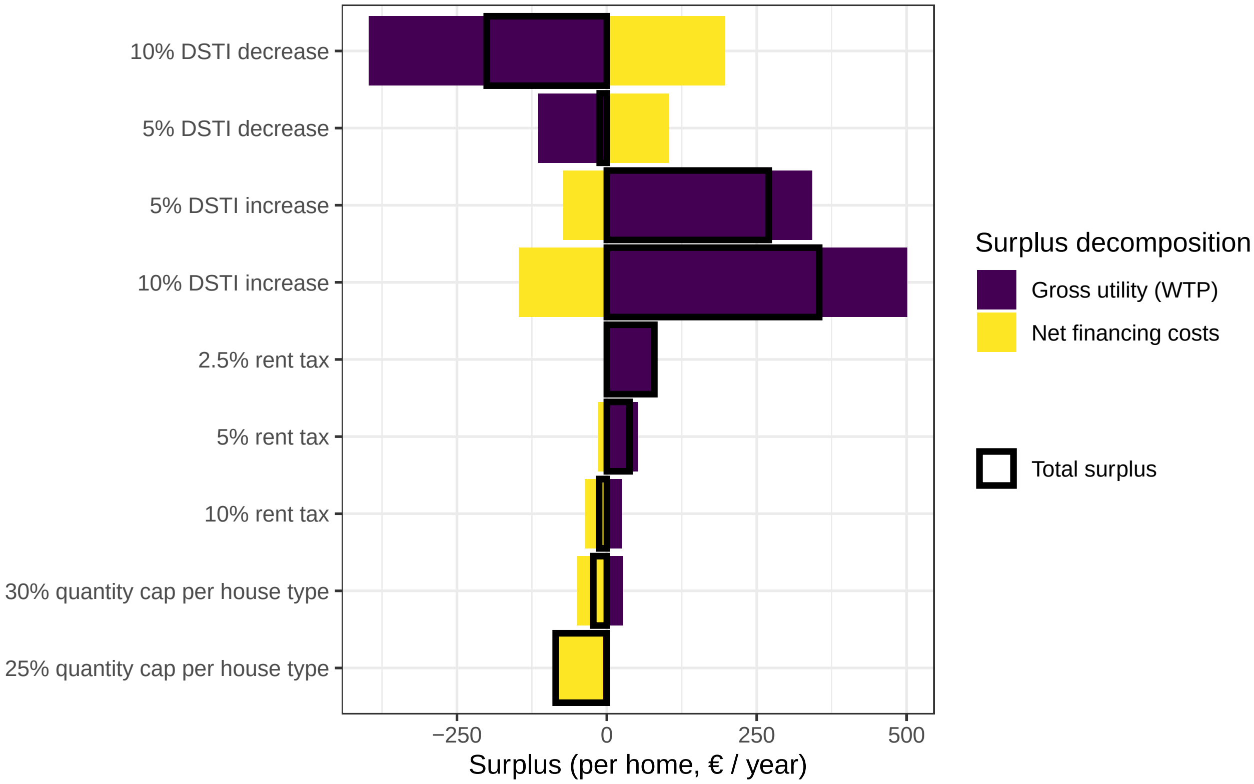

Note: The figure shows the effect on total surplus of various policy measures in 2019. Changes to DSTI measure a uniform increase or decrease of the given percentage. Rent taxes measure the given uniforn tax on rental income in addition to the taxes that exist in the counterfactual. Quantity caps measure the effect of a constraint on the fraction of rentals per house type. Gross utility measures changes in total willingness to pay, i.e. we have deducted transfers between households, investors and the government. Net financing costs measures changes in the financing cost of houses. For houses owned by households, financing costs are given by the user cost of housing. For houses owned by investors, financing costs are given by the cost of capital. Total surplus measures the change to total surplus, which is the sum of gross utility and net financing costs.

these effects into two effects. First, we compute households' gross consumer surplus. This equals the willingness to pay for their predicted choice, reduced by our estimate of fiscal subsidies. Changes to gross utility come from two sources of allocative efficiency. The first is the extent to which households' choices match with their tenure preferences, i.e. whether they buy or rent (the "tenure choice channel"). The second is the extent to which their choices match with their house preferences, i.e. which house they live in. (the "house choice channel"). We measure the total financing cost of all houses, as given by the opportunity cost of capital. For households, this equals the user cost of housing, as we explain in Section 4. For investors, this is the estimated hurdle rate of capital. There are hence possible welfare gains (losses) when a house is sold to an agent with a lower (higher) financing cost.

Small changes to DSTI limits can have sizeable impacts on total surplus in the housing market. For example, if we uniformly increase DSTI limits with 5%, total surplus increases with €270 per household per year, or over €1.3 billion per year in total. The other measures we consider can increase total surplus, but not to the same extent as relaxing DSTI limits. However, contrary to relaxing DSTI limits, taxes and quantity caps have no negative effects on financial stability.

Figure 11: Waterbed effects: the effects of a 30% quantity cap on rentals per house type

(a) All houses (b) Only unaffected houses

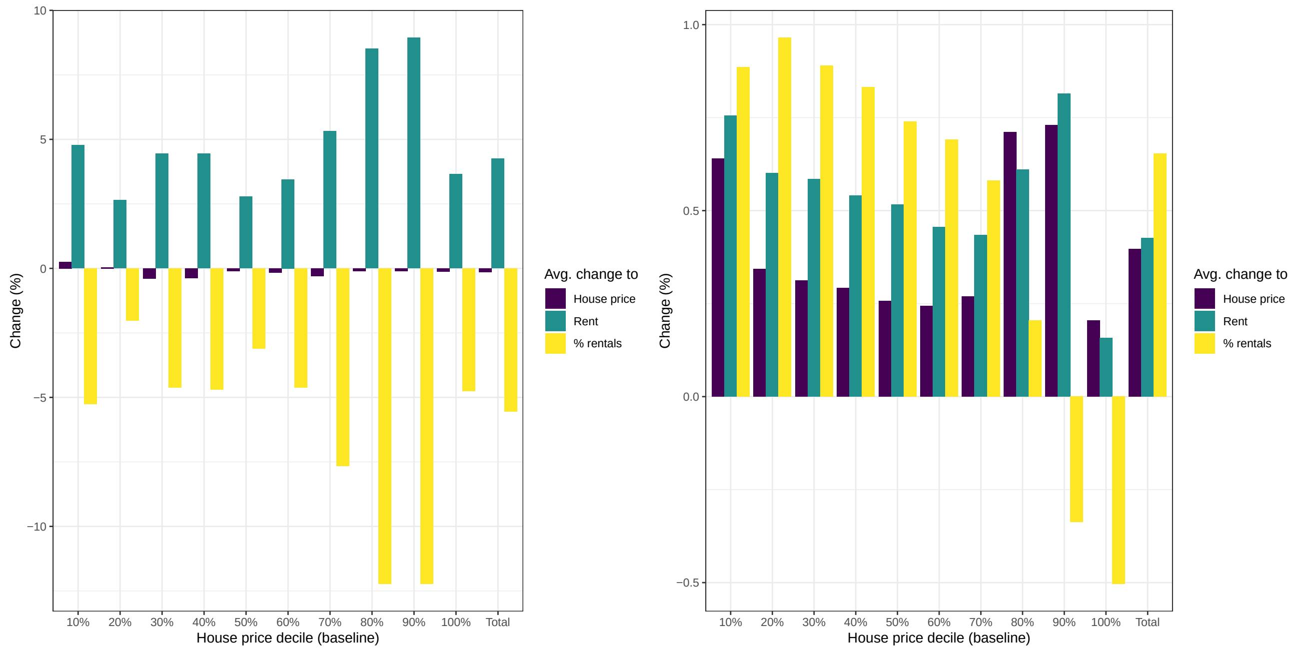

Note: The figures shows the effect on house prices, rents and fraction rentals of a 30% cap on rentals, for 2019. The left panel shows the effects on all house types, where the right panel shows only the effects on houses that are unaffected by this policy as they have less than 30% rentals in the baseline.

Taxes and quantity caps do not increase total surplus as much for the following reasons. The first is that, as they do not change the choice sets of households that already own, they only reduce the deadweight loss of DSTI limits through the tenure choice channel. The second is that uniform measures risk overshooting for some housing types. If we again look at Figure 7, we need a cap on the fraction of rentals of about 10% to counteract the effect of DSTI limits for the middle housing type. Such a cap however reduces the proportion of rentals to the left of the constrained efficient amount in the left-most submarket. To prevent undershooting, a low quantity cap or tax is required. But this limits the effect of such measures for housing types whose occupants are most restricted. The third reason is that such measures can introduce waterbed effects. This can be seen in Figure 11, which shows the effects of a 30% quantity cap. As the left panel shows, this leads to a 5% reduction in the number of rentals. Rents increase due to lower supply, while owner-occupied dwellings are on average unaffected. The right panel shows the same effects for house types that are not affected by this cap, i.e. that had fewer than 30% rentals in the baseline. We see that the proportion of rentals increase by just under .7%. These house types hence move away from the constrained efficient equilibrium as demand for rentals shifts to these unconstrained house types. Welfare gains for affected housing types are thus partially offset by welfare losses for unaffected housing types. Increases to DSTI limits cannot cause overshooting or waterbed effects: demand always shifts towards the no-DSTI demand curve for every housing type.

Finally, we note that there is no one-to-one relationship between a policy's effect on the number of constrained renting households and its effect on total surplus. This is for two reasons. First, the number of constrained renting households is only indicative of the size of the tenure choice channel but not of the house choice channel. Since the number of constrained renters moves relatively weakly with changes in DSTI limits, it must be the case that the house choice channel is an important source of deadweight loss. Hence, the number of constrained renters is not a sufficient statistics for the total deadweight loss of DSTI limits. The house choice channel is not reflected in the number of constrained renters. The second reason is that policies can introduce additional distortions. We saw examples of the latter mechanism above in the form of overshooting and waterbed effects.

5.4 Mechanisms behind DSTI limits

To gain a better understanding of the mechanisms of DSTI limits, we zoom in on the effects of a 5% increase. We find qualitatively similar mechanisms in the other counterfactuals in which we vary DSTI limits. An increase in DSTI limits leads to both higher house prices and higher rents (Figure 12a). Increasing house prices are expected, as households are less constrained and the demand for owner-occupied housing shifts to the right. Perhaps more surprising is that rents increase as well, because a right-ward shift in the demand for owner-occupied housing must mean a left-ward shift in the demand for rentals. There are, however, two countervailing forces. The first is that, as the price for owner-occupied housing increases, the demand for rentals again increases as owner-occupation and renting are substitutes. The second is that buy-to-let investors sell some of their stock, which constrains the supply of rentals and hence pushes up the price. In our setting, it turns out that that the latter two upward forces on rents are larger than the first downward force.

This relaxation of DSTI limits increases both consumer and investor surplus. Households are on average €342 per year better off (Figure 12b. This is because the utility they obtain in the housing market, as measured by their gross consumer surplus, goes up. This reflects two forces. First, there are more gains from trade as households are less constrained. Hence, previously constrained households with a higher willingness to pay for housing quality can move up the housing ladder by trading with households that have a lower willingness to pay for housing quality. Second, previously constrained renting households that prefer to own are now able to do so. All other effects measure transfers between households, investors and the government. As can be seen in the figure, these effects essentially cancel out on aggregate. Investors are somewhat better off, with an increase in surplus of €34 per house on average. They benefit because rents increase more than the amount of rentals decreases, so that total rental income goes up. Moreover, the prices of their properties increases.

5.5 Redistributive effects of DSTI limits and other policies

We now disaggregrate the effects on consumer surplus. Again focussing on the 5% DSTI increase case, we do not find any evidence for redistributive effects. As Figure 12c shows, households across the income distribution benefit. However, all increases in surplus that result from a relaxation of DSTI limits accrue to current property owners: the consumer surplus of households that rent or are in the outside option in our baseline does not change (Figure 12d). This result follows from our assumption that the housing stock is fixed. If, for

Figure 12: Equilibrium effects of a 5% increase in DSTI limits

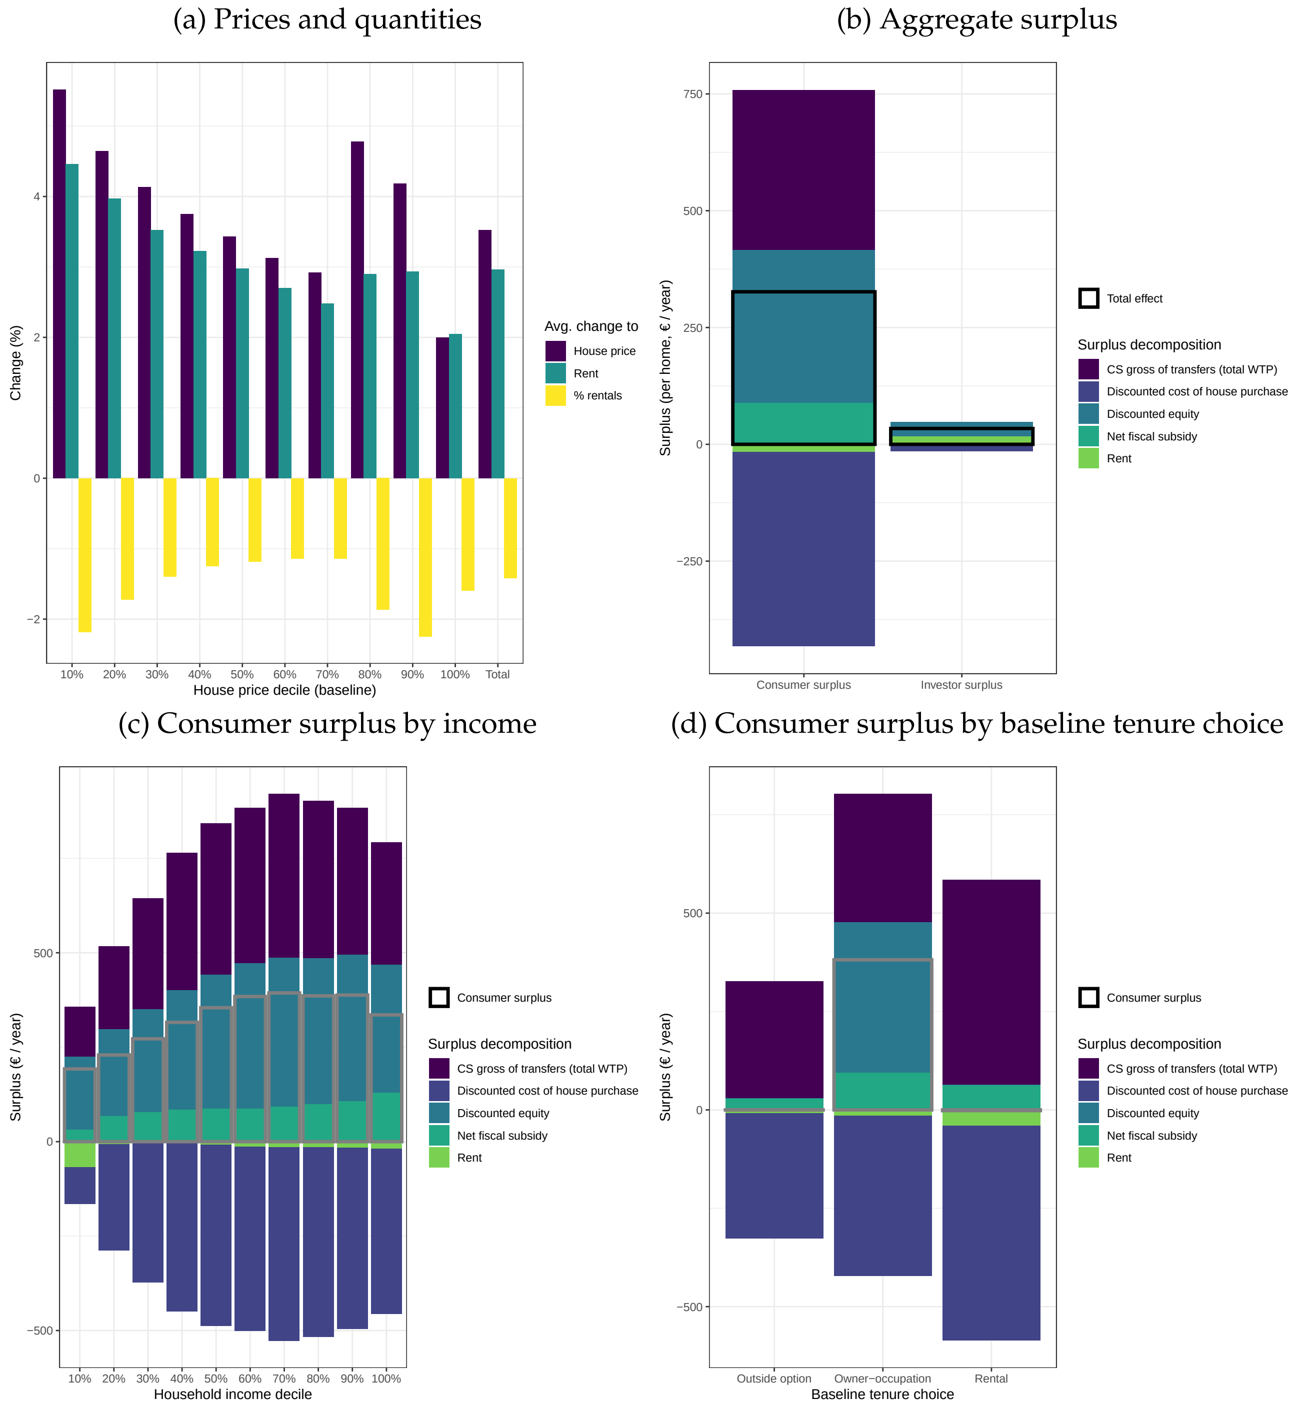

Note: The figures show the effects of a uniform increase of 5% of DSTI limits. Panel (a) shows predicted effects on house prices, rents and fraction of rentals, by government-assed house value. Panel (b) shows the effects on consumer and investor surplus, (c) decomposes consumer surplus by income decile, and (d) by baseline tenure choice. See main text for decomposition of consumer surplus.

Figure 13: Equilibrium effects of a 2.5% tax on rental income

(a) Prices and quantities (b) Consumer surplus by baseline tenure choice

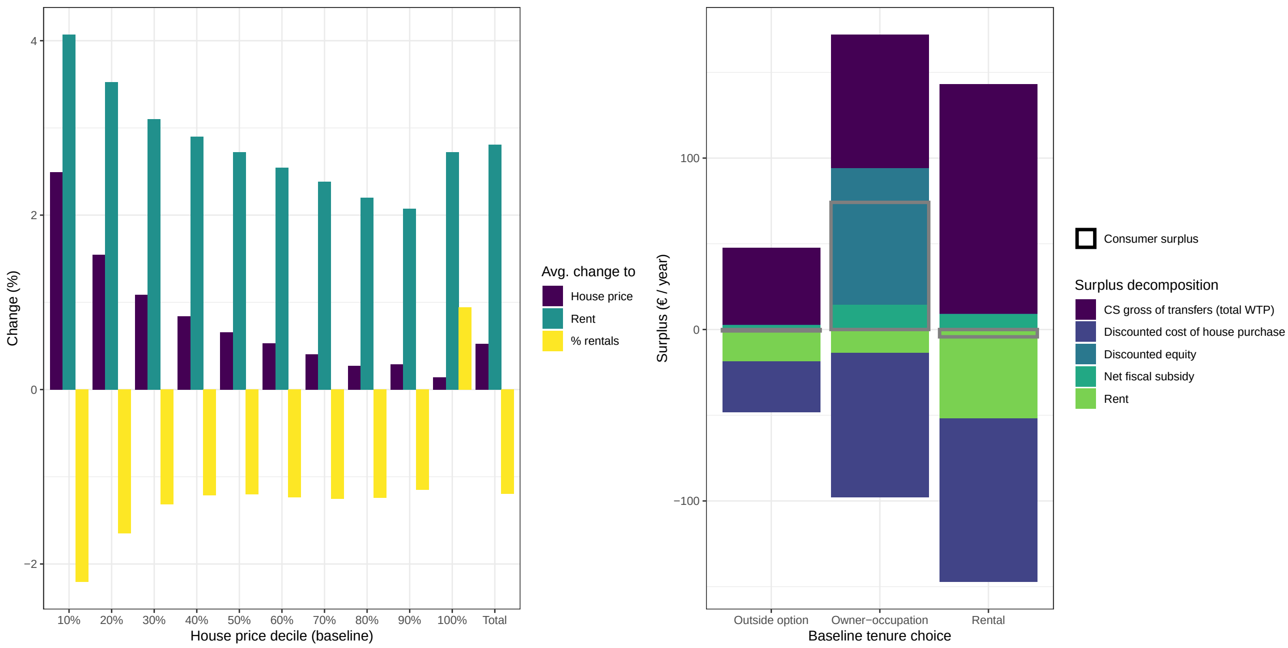

example, households' demand shifts outwards as a result of a relaxation of DSTI limits, the equilibrium price then increases one-to-one with households' higher willingness to pay. In other words, while all welfare effects in our setting occur due to shifts in the demand curve, all changes in welfare are distributed to the supply side. In our case, the supply side consists of households and investors that already own properties. Hence, while relaxing DSTI limits benefits households broadly, it may do little for housing market outsiders when supply is inelastic. However, in a more general model with housing construction, these households may be better off as increases in demand will lead to more supply instead of price increases.

Since changes to DSTI limits do not help housing market outsiders, a natural question is whether the other policies we consider can do so. Here, we focus on a 2.5% rent tax, as this is the policy that we find gives the greatest increase in total surplus across the non-DSTI policies we consider. Such a tax does not directly help housing market outsiders: households that rent or are in the outside option in the baseline are slightly worse off (Figure 13b). The mechanism here is similar as when changing DSTI limits: since total housing supply is inelastic, increases in aggregate willingness to pay go to property owners. However, the government does raise some tax revenues in this case. If the government has access to lump sum transfers, it could use this revenue to make outsiders better off.

6 Discussion

In this paper, we discussed the effects of macroprudential policies on the housing market. We focused on DSTI limits, which limit the households' mortgage relative to their income.

Exploiting changes to DSTI limits in the Netherlands, we showed that DSTI limits have a significant impact on household tenure choice: the elasticity of the probability to rent rather than own with respect to the maximum mortgage is -8. In equilibrium, this means DSTI limits distort tenure choice towards renting. We built an equilibrium model of the housing market with buy-to-let investors to quantitfy this effect and its welfare consequences.

While we have focused on DSTI limits, similar mechanisms probably arise for other policies that constrain mortgage borrowing, such as Loan-to-Value limits. From the point of the household it does not matter what constrains its maximum mortgage, so that any constraint that limits borrowing likely influences tenure choice. This gives buy-to-let investors an incentive to increase their investment into the housing market, increasing the number of rentals. However, the quantitative effects may well be different than for DSTI limits. It would be interesting to extend our results to other macroprudential policies to see if there are differences in their effects on housing market equilibrium.

Our results imply that policy makers face a trade off between the macroprudential goals of DSTI limits and deadweight losses in the housing market. However, there is no one-to-one relationship between the tightness of DSTI limits and this deadweight loss. For example, increasing DSTI limits by 5% gives almost the same increase in total surplus as increasing them by 10%, which reflects that many households in the Netherlands are on the margin of being constrained. Hence, to make the right trade-off requires precise knowledge of the macroprudential effects of DSTI limits. While there is a growing body of research on the effects of macroprudential policies on financial stability, bank risk, etcetera (Galati and Moessner, 2013; Claessens, 2015), these typically do quantify the welfare gains of these policies. This makes it difficult to weigh them against their effects on the housing market. We view this as an important area for future research.

References

Armstrong, Timothy B., "Large Market Asymptotics for Differentiated Product Demand Estimators With Economic Models of Supply," *Econometrica*, 2016, *84* (5), 1961–1980. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA10600.

Bajari, Patrick, C. Lanier Benkard, and John Krainer, "House prices and consumer welfare," *Journal of Urban Economics*, November 2005, *58* (3), 474–487.

Bayer, Patrick, Fernando Ferreira, and Robert McMillan, "A Unified Framework for Measuring Preferences for Schools and Neighborhoods," *Journal of Political Economy*, 2007, *115* (4), 588–638. Publisher: The University of Chicago Press.

—, Robert McMillan, Alvin Murphy, and Christopher Timmins, "A Dynamic Model of Demand for Houses and Neighborhoods," *Econometrica*, 2016, *84* (3), 893–942. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA10170.

—, —, and Kim Rueben, "An Equilibrium Model of Sorting in an Urban Housing Market," Technical Report w10865, National Bureau of Economic Research November 2004.

Bekkum, Sjoerd Van, Marc Gabarro, Rustom M Irani, and Jose-Luis Peydro, "Take It to the Limit? The Effects of Household Leverage Caps," 2019.

Bracke, Philippe, "House Prices and Rents: Microevidence from a Matched Data Set in Central London," *Real Estate Economics*, June 2015, *43* (2), 403–431.

Campbell, Sean D., Morris A. Davis, Joshua Gallin, and Robert F. Martin, "What moves housing markets: A variance decomposition of the rent–price ratio," *Journal of Urban Economics*, September 2009, *66* (2), 90–102.

Case, Karl E. and Robert J. Shiller, "The Efficiency of the Market for Single-Family Homes," *The American Economic Review*, 1989, *79* (1), 125–137. Publisher: American Economic Association.

Chetverikov, Denis and Daniel Wilhelm, "Nonparametric Instrumental Variable Estimation Under Monotonicity," *Econometrica*, 2017, *85* (4), 1303–1320. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA13639.

__, Dongwoo Kim, and Daniel Wilhelm, "Nonparametric Instrumental-Variable Estimation," *The Stata Journal: Promoting communications on statistics and Stata*, December 2018, *18* (4), 937–950.

Claessens, Stijn, "An Overview of Macroprudential Policy Tools," *Annual Review of Financial Economics*, 2015, *7* (1), 397–422. _eprint: https://doi.org/10.1146/annurev-financial-111914-041807.

Dougherty, Ann and Robert Van Order, "Inflation, Housing Costs, and the Consumer Price Index," *The American Economic Review*, 1982, *72* (1), 154–164. Publisher: American Economic Association.

Ewijk, Casper Van and Arjan Lejour, "De lage rente biedt kans om de fiscale behandeling van de eigen woning te hervormen," *Economisch Statistische Berichten*, 2017, pp. 46–59.

Favilukis, Jack, Sydney C. Ludvigson, and Stijn Van Nieuwerburgh, "The Macroeconomic Effects of Housing Wealth, Housing Finance, and Limited Risk Sharing in General Equilibrium," *Journal of Political Economy*, February 2017, *125* (1), 140–223. Publisher: The University of Chicago Press.

Galati, Gabriele and Richhild Moessner, "Macroprudential Policy – a Literature Review," *Journal of Economic Surveys*, 2013, *27* (5), 846–878. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1467-6419.2012.00729.x.

Glaeser, Edward L. and Erzo F. P. Luttmer, "The Misallocation of Housing Under Rent Control," *American Economic Review*, September 2003, *93* (4), 1027–1046.

__ and Joseph Gyourko, "Arbitrage in Housing Markets," Technical Report w13704, National Bureau of Economic Research December 2007.

__, Joshua D. Gottlieb, and Joseph Gyourko, *7. Can Cheap Credit Explain the Housing Boom?*, University of Chicago Press, 2013.

Halket, Jonathan, Lars Nesheim, and Florian Oswald, "The Housing Stock, Housing Prices, and User Costs: The Roles of Location, Structure, and Unobserved Quality," *International Economic Review*, 2020, *61* (4), 1777–1814. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/iere.12475.

Han, Lu, L Rachel Ngai, and Kevin D Sheedy, "To Own or to Rent?The Effects of Transaction Taxes on Housing Markets," 2022.

Head, Allen and Huw Lloyd-Ellis, "Housing Liquidity, Mobility, and the Labour Market," *The Review of Economic Studies*, October 2012, *79* (4), 1559–1589.

Henderson, J. V. and Y. M. Ioannides, "A Model of Housing Tenure Choice," *The American Economic Review*, 1983, *73* (1), 98–113. Publisher: American Economic Association.

Hill, Robert J. and Iqbal A. Syed, "Hedonic price–rent ratios, user cost, and departures from equilibrium in the housing market," *Regional Science and Urban Economics*, January 2016, *56*, 60–72.

Ke, Guolin, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, and Tie-Yan Liu, "LightGBM: A Highly Efficient Gradient Boosting Decision Tree," in "Advances in Neural Information Processing Systems," Vol. 30 Curran Associates, Inc. 2017.

Kiyotaki, Nobuhiro, Alexander Michaelides, and Kalin Nikolov, "Winners and Losers in Housing Markets," *Journal of Money, Credit and Banking*, 2011, *43* (2-3), 255–296. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1538-4616.2011.00374.x.

Lambie-Hanson, Lauren, Wenli Li, and Michael Slonkosky, "Real estate investors and the U.S. housing recovery," *Real Estate Economics*, 2022, *50* (6), 1425–1461. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/1540-6229.12396.

Linneman, Peter, "An economic analysis of the homeownership decision," *Journal of Urban Economics*, March 1985, *17* (2), 230–246.

McFadden, Daniel, "Computing willingness–to–pay in random utility models," in "Trade, theory and econometrics," Routledge, 2012, pp. 275–296.

Michielsen, Thomas, Stefan Groot, and Raoul Van Maarseveen, "Prijselasticiteit van het woningaanbod," Technical Report, Centraal Planbureau 2017.

Mills, James, Raven Molloy, and Rebecca Zarutskie, "Large-Scale Buy-to-Rent Investors in the Single-Family Housing Market: The Emergence of a New Asset Class," *Real Estate Economics*, 2019, *47* (2), 399–430. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/1540-6229.12189.

Nevo, Aviv, "Measuring Market Power in the Ready-to-Eat Cereal Industry," *Econometrica*, 2001, *69* (2), 307–342. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/1468-0262.00194.

Olden, Andreas and Jarle Møen, "The triple difference estimator," *The Econometrics Journal*, September 2022, *25* (3), 531–553.

Poterba, James M., "Tax Subsidies to Owner-Occupied Housing: An Asset-Market Approach," *The Quarterly Journal of Economics*, 1984, *99* (4), 729–752.

Rouwendal, Jan, Florian Sniekers, and Jia Ning, "Borrowing constraints, housing tenure choice and buy-to-let investors: An assignment model," 2023.

Sculley, D., "Web-scale k-means clustering," in "Proceedings of the 19th international conference on World wide web" WWW '10 Association for Computing Machinery New York, NY, USA April 2010, pp. 1177–1178.

Small, Kenneth A. and Harvey S. Rosen, "Applied Welfare Economics with Discrete Choice Models," *Econometrica*, 1981, *49* (1), 105–130. Publisher: [Wiley, Econometric Society].

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

Statistics Netherlands, "Bijna half miljoen woningen in particuliere verhuur," March 2019. Last Modified: 2019-04-01T00:00:00+02:00.

__ , "Aandeel woonlasten in inkomen gedaald in 2021," June 2022. Last Modified: 2022-06-09T10:00:00+02:00.

Train, Kenneth E., *Discrete choice methods with simulation*, Cambridge university press, 2009.

Tzur-Ilan, Nitzan, "Adjusting to Macroprudential Policies: Loan-to-Value Limits and Housing Choice," *The Review of Financial Studies*, May 2023, p. hhad035.

Table A1: Datasets used in the analysis

Dataset	Content
Microdata:	
Gbapersoontab	Demographics of individuals
Gbahuishoudensbus	Household information
Inhatab	Household income
Vehtab	Household wealth
GBAADRESOBJECTBUS	Addresses of individuals
Vslcoordtab	Coordinates of addresses and buildings
Eigendomwozbagtab	House value
Eigendomtab	Type of property and owner
Levcyclwoonnietwoonbus	Life-cycle of houses
Energieverbruiktab	Energy and gas usage
Bestaandekoopwoningen	Cadastre data with house sale prices
Woonbase	Combination of housing microdata and survey data
Survey data:	
Huurenquete	Rent survey
Woon	Housing survey
Other sources	
Nibud Financieringslastnormen	DSTI norms

Supplementary Material

Appendix A Data set description

In this appendix we describe the data sets used in the paper. Table **??** lists all sources, for which we provide additional information in the following section.

A.1 Datasets

1. Gbapersoontab contains register data with demographic information about every person registered at a Dutch muncipality.
2. Gbahuishoudensbus contains register data with information about the household that a person belongs to, their position in the household, and start and end date of the household. Allows linking of persons to households.
3. Inhatab contains register data about the incomes of all households.
4. Vehtab contains register data about the wealth of all households.
5. GBAADRESOBJECTBUS contains register data with address information for all persons. Together with Gbahuishoudensbus, this allows the linking of houses, persons, and households.
6. Vslcoordtab contains coordinates for each house in the Netherlands. This allows for spatial clustering of homes.

7. Eigendomwozbagtab contains register data with property values of all houses in the Netherlands.
8. Eigendomtab contains register data with the type of property (rental or owner-occupied) and type of owner (Own house, Housing association, Other landlord).
9. Levcyclwoonnietwoonbus contains register data about the life-cycle of houses with information about usage, size, date of modification, and year of construction of houses.
10. Energieverbruiktab contains register data on energy and gas usage of houses.
11. Bestaandekoopwoningen contains register data with sales prices for all house sales as registered with the Cadastre. Importantly, this excludes the sales price of newly built apartments.
12. Huurenquete contains survey information about rent and service costs of rental housing in the Netherlands.
13. Woon contains survey information about the housing situation and preferences of respondents. The survey responses are enriched with register data. Other than the other datasets, this is only measured triennial.
14. Woonbase contains register data about housing, households, and persons similar to above mentioned databases. This dataset is only available for 2019. As this dataset contains more rental data for social housing than the Huurenquete, we use it to improve the rent imputations.
15. Nibud Financieringslastnormen contains the DSTI norms for all years, income, and interest rate brackets.

Appendix B Imputation of rents

We lack register data on two important housing variables: whether a rental house is market-rent or social, and the rent. We therefore impute these based on survey data, which we described in the previous section.

We impute the market-rent status for 2019 and back date these to the earlier years of our analysis, i.e. we assume that a house that was market-rent in 2019 also was market-rent before 2019. We use a boosted gradient forest (Ke et al., 2017) to compute a predictive model whether a rental is social or market-rate. In Figure B1a, we measure the goodness-of-fit of this model on a hold-out sample. Our model has an Area Under Curve of .9968, which means that are able to predict well whether a given house is market-rent. The most important predictive variables are i) whether the house is included in the Woonbase data set (see above), as this data set includes all social rentals but only a subset of market-rate rentals, ii) government-assed house value, iii) location dummies. Our model generates a probability for every rental that is market-rate. We pick a cut-off value for this probability as follows. We know the aggregate fraction of rental houses from a survey on rentals (Huurenquete). We pick the cut-off value such that the aggregate fraction in our sample

Figure B1: Quality of imputations

(a) Liberalization status (b) Rent

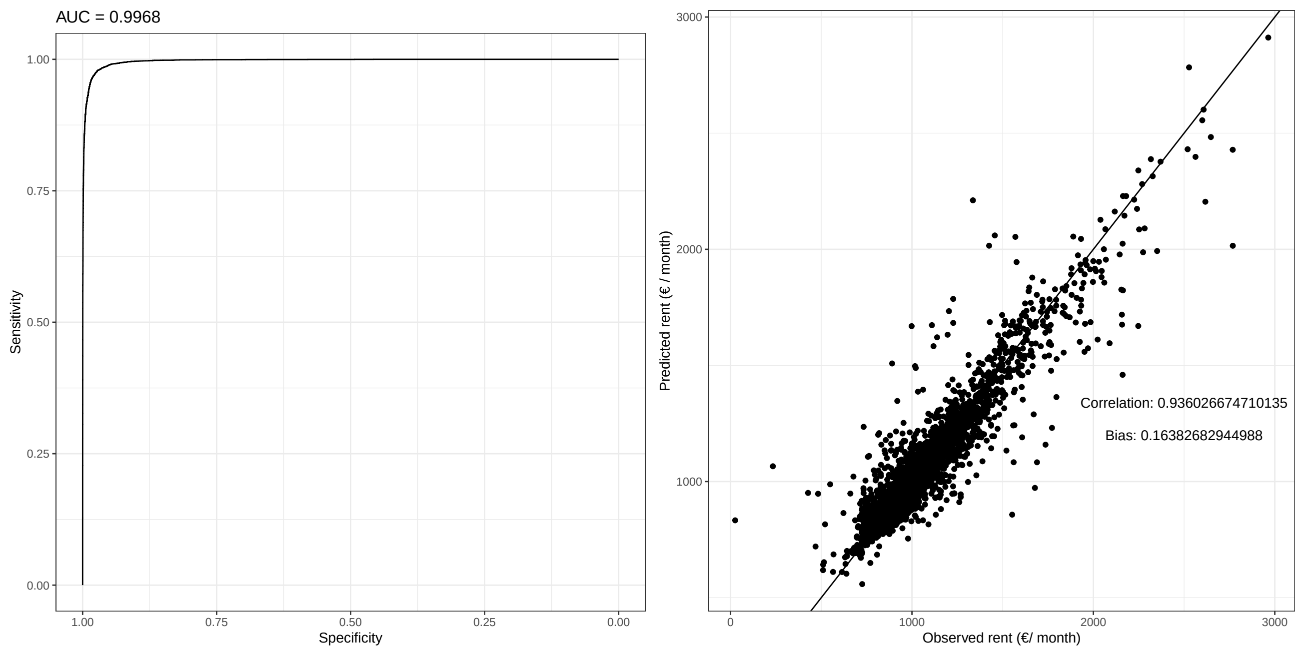

Note: The figures shows the fit of our imputation procedures for the liberalization and rent of houses in the rental market. Goodness-of-fit is computed on a 10% hold-out sample in both cases.

matches this known fraction. We note that for most rentals, predicted probabilities are below 1% or above 95%, so that only a few houses are impacted by the choice of cut-off.

To impute rents, we proceed as follows. On a sample of known market-rate rentals, we train a boosted gradient forest to predict the rent. We do this separately for every year in our sample. The most important variables in our model are i) government-assessed house value, ii) location dummies, iii) household income of the occupant. Figure B1b shows the predictive power of this model on a hold-out sample for 2019; other years are similar. As the figure shows, the model is able to predict rents very well. Important for our application is that we have low bias: in our model, we take as rent of a house type the average over multiple houses. A higher variance is averaged away, while a bias would persist. We have a bias under €1 per month.

Appendix C Clustering of houses

As we describe in the main text, we use batched K-means clustering to divide all houses in our sample into housing types, which contain houses with similar characteristics. We have 5,865 housing types in our sample. We cluster on i) location (x- and y-coordinate), ii) government-assessed house value, iii) square footage.

Figure C1 shows the house types in our sample. On the left, we see that most house types contain a limited number of houses. The average number of houses per house type is 55.8. There are house types that contain only a single house. The right shows the distribution of the centers of the estimated clusters: if two house types have similar government-assessed values and square footages, a house is assigned to the closest house

Figure C1: House types

(a) Number of houses per house type

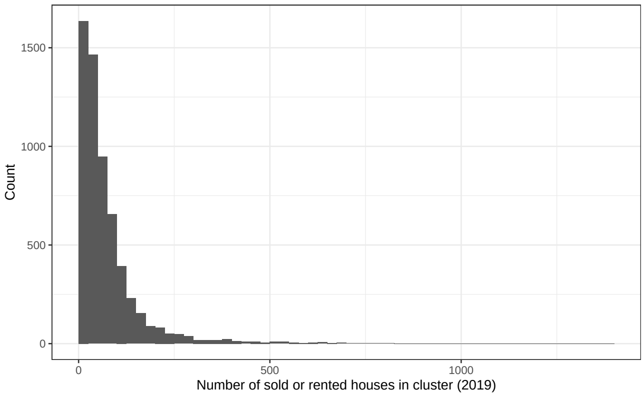

(b) Location of house types

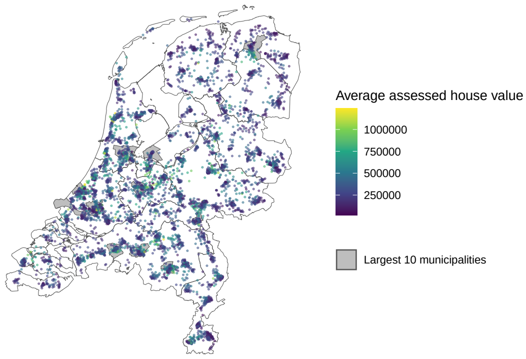

type. House types are centered around larger cities. There are more house types in the more densely distributed West of the country.

Figure C2 shows the homegeneity of our house types. It shows the distribution of the coefficient of variation (CV), i.e. the ratio of the mean and the standard deviation. Panels (a) and (b) show the coefficient of variation for variables for which the algorithm tries to minimize the variation within clusters. Indeed, for most clusters we find a small CV, with most clusters having a CV under .1. Panels (c) through (f) repeat this exercise for characteristics that we have not instructed the algorithm to cluster on. They therefore show if the house types are also homogeneous for characteristics that are not homogeneous "by construction". CV's are similar as in panels (a) and (b), with the exception of the fraction of apartments. However, the CV should be interpreted a bit differently for binary variables. For a binary variable, the CV equals $\sqrt{p/(1-p)}$, where p is the fraction of successes. To obtain a CV under .1, comparable with the other characteristics, then requires that over 99% of houses in house type is an apartment (or not). A CV of .25, which is what we have for most house types, means that over 94% of houses is an apartment (or not). Hence, house types are also very homogeneous on this dimension.

Appendix D Additional model estimates

We present some additional results from our model. We start with demand-side estimates. Table D1 shows the heterogeneity in willingness to pay for housing attributes. There is indeed significant heterogeneity between households; interaction terms all have the intuititive sign. Table D2 shows second-stage estimated coefficients, in which we regress house type mean utility on prices and house type-year fixed effects. Column (ii) shows that the price sensitivity towards owner-occupied houses increases with the mortgage interest rate, as expected. Even though this effect is not estimated very precisely, we use this as our main specification. In columns (iii) and (iv) we do not control for house type-year fixed effects, i.e. we just regress mean utility on prices. We find that without our correction, we would estimate households to be less price sensitive. Columns (v) and (vi) additionally control for Rent and Rent × House type fixed effects, respectively. That is, we allow for

Figure C2: Homogeneity of house types

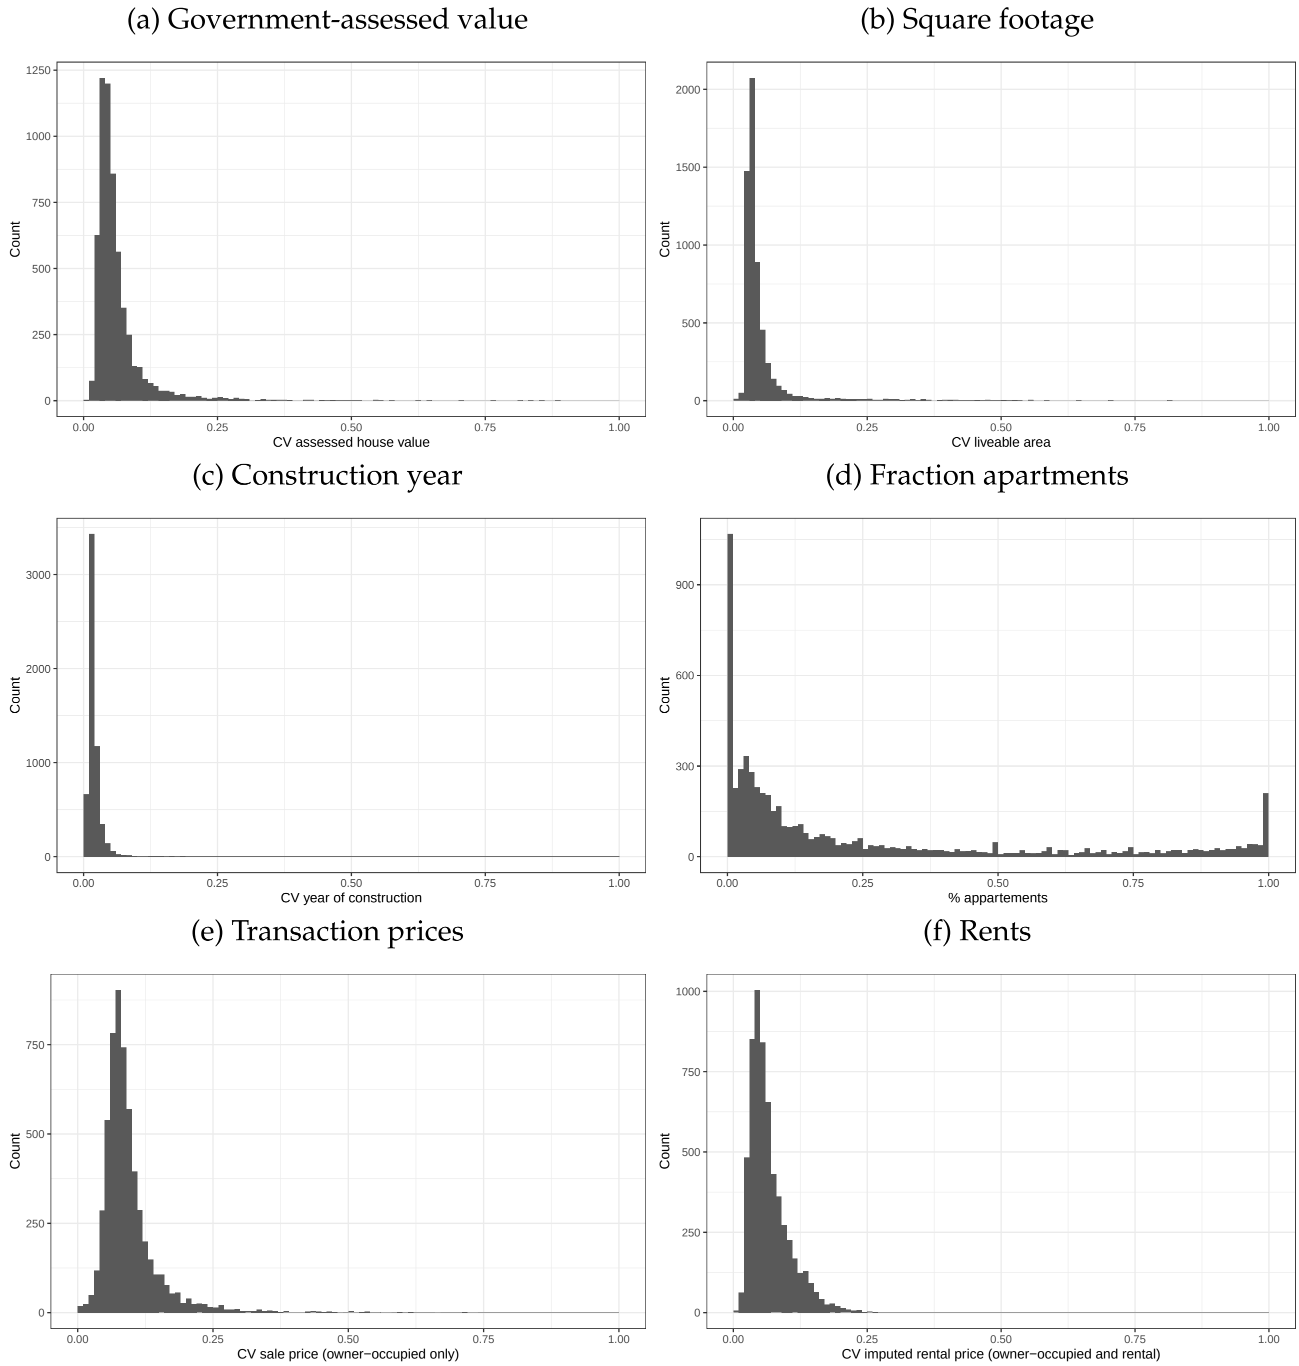

Note: The figures show the distribution of the coefficient of variation (CV) of housing characteristics amongst house types. The coefficient of variation is the mean divided by the standard deviation. House types are computed using *K*-means clustering, see the main text. Government-assessed value and square footage are variables that are directly targeted in clustering algorithm. Construction year, the fraction of apartments and prices are not.

Table D1: Heterogeneity in willingness to pay for housing attributes

	Size (m^2)	Urbanity	Apartement
Avg. age of adults in household (years)	0.17	-0.26	-34.88
No. people in household	1.66	16.78	-965.51
Yearly household income (log €)	9.43	223.18	-480.14
Household wealth (€ 10000)	0.06	-0.80	-4.91

Note: The Table shows the effect of household characteristics in the willingness to pay (WTP) for housing attributes in terms of yearly rent. Every cell shows the marginal effect of an increase in the row variable on the willingness to pay for the column variable. Urbanity measures the urban density of the housing type on a scale 1–5, with higher being more urban. Apartment is a dummy containing whether the house is an apartment (1) or a single-family dwelling (0). WTP is calculated by dividing the interaction parameter between household and house characteristics, γ^{β}, by the mean coefficient on rent, μ^{α^R}.

structural quality differences between rentals and owner-occupied houses. As the results show, we get large standard errors on the price coefficients in these specifications.[46] We conclude that we are not able to precisely differentiate tenure preferences from quality differences in our setup.

We continue with supply-side estimates. Table D3 shows the linear equivalent of the non-parametric instrumental variable (NPIV) approach that we take. We present this case as the diagnostics for weak variables are better understood. We have large first-stage F-statistics and significant second-stage coefficients. Figure D1 shows our the fit NPIV regression visually. We compare it to a linear IV estimator, as well as regular linear and non-linear regressions. Both for the linear IV and the NPIV, we find a steeper supply curve than in the corresponding models that do not instrument for the rent/price ratio.

[46]The fact that some price coefficients become positive is not itself problematic, as what matters is the sign of the mean + idiosyncratic price coefficient. In these specifications, the majority of households still have negative coefficients on price.

Table D2: Estimated demand model coefficients

Dependent Variable:	Mean utility					
	(i)	(ii)	(iii)	(iv)	(v)	(vi)
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Variables						
Purchase price	$-1.22 \times 10^{-6***}$	$-9.65 \times 10^{-7***}$	$5.22 \times 10^{-7**}$	6.23×10^{-8}	2.66×10^{-7}	2.24×10^{-7}
	(8.94×10^{-8})	(2.25×10^{-7})	(1.32×10^{-7})	(1×10^{-5})	(2.35×10^{-7})	(1.57×10^{-7})
Yearly rent	-0.0010***	-0.0010***	-0.0004***	-0.0004***	0.0001	-3.36×10^{-5}
	(2.88×10^{-5})	(2.95×10^{-5})	(3.99×10^{-5})	(4.04×10^{-5})	(7.54×10^{-5})	(6.34×10^{-5})
Purchase price $\times$ mortgage interest rate		-1.23×10^{-7}		2.06×10^{-7}	-6.58×10^{-9}	5.13×10^{-8}
		(9.9×10^{-8})		(1×10^{-5})	(9.7×10^{-8})	(8.32×10^{-8})
(Intercept)			-5.808***	-5.812***		
			(0.0452)	(0.0457)		
Fixed-effects						
House-Year	Yes	Yes			Yes	Yes
Rental					Yes	
House-Rental						Yes
Fit statistics						
Observations	58,650	58,650	58,650	58,650	58,650	58,650
R^2	0.88511	0.88512	0.01903	0.01906	0.88609	0.96813
Within R^2	0.15154	0.15159			0.00013	0.00020

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Note: The Table shows coefficients from the second-stage demand estimation, in which mean utility is regressed on price. The models in the first to columns contain house type-year fixed effects, which allow us to identify the effect of price on utility. The middle two columns omit these fixed effects, for comparison. The latter two columns also control for rent or rent/house type fixed effects, to control for possible quality differences between rentals and owner-occupied houses.

Table D3: Coefficients of linear version of supply side model

	2015		2016		2017		2018		2019	
	% rentals	Rent/price	% rentals	Rent/price	% rentals	Rent/price	% rentals	Rent/price	% rentals	Rent/price
Variables										
(Intercept)	-0.1357***	0.0314***	-0.1181***	0.0319***	-0.1156***	0.0351***	-0.1482***	0.0372***	-0.1462***	0.0367***
	(0.0068)	(0.0010)	(0.0061)	(0.0008)	(0.0063)	(0.0006)	(0.0062)	(0.0004)	(0.0061)	(0.0003)
User cost	9.447***		9.232***		9.223***		11.25***		11.99***	
	(0.3262)		(0.3147)		(0.3447)		(0.3473)		(0.3590)	
Frac. rentals		0.2945***		0.2476***		0.2008***		0.1208***		0.0938***
		(0.0157)		(0.0128)		(0.0116)		(0.0065)		(0.0049)
Fit statistics										
Observations	5,865	5,865	5,865	5,865	5,865	5,865	5,865	5,865	5,865	5,865
F-test	2,096.4	2,780.1	2,040.4	2,851.9	2,570.9	1,907.1	3,645.0	1,239.2	3,927.8	1,169.4
Kleibergen-Paap		826.97		717.79		782.85		1,063.3		1,115.6

Heteroskedasticity-robust standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Note: The Table shows first-stage and second-stage coefficients of linear instrumental variable regressions of rent/price ratios on the fraction of rental, year by year. The instrument is the average user cost per housing type as computed by our demand model when houses are only horizontally differentiated.

D.1 Counterfactuals with quality differences

In this Appendix, we re-compute some of our main counterfactuals using a demand specification that allows for structural quality differences between rentals and owner-occupied houses. In particular, we use specification (vi) from Table D2, which allows for a different quality difference for every house type. We note that changing to this specification has two effects. Of course, the demand function is different so that different equilibrium we (in general) obtain different equilibrium prices in our counterfactuals. But since we also use the user cost of housing, which is higher in specification (vi), to convert changes in house prices to yearly money-metric utility, we also obtain different surplus measures even for the same equilibrium prices.

In Figure D2a we show the number of (constrained) renters in our counterfactuals. Comparison with Figure Figure 9 shows a similar picture, although the changes to the number of renters are generally a bit smaller than in the main text.

Figure D2b shows changes to total surplus for our counterfactuals. Compared to our main results, we find similar effects of changes to DSTI limits, but typically slightly more positive effects of the other policies we consider. This can largely be attributed to the fact that the house price increases that occur in all these counterfactuals receive a larger welfare weight because the average user cost is higher using this demand function.

Figure D1: Distribution of investors' hurdle rates

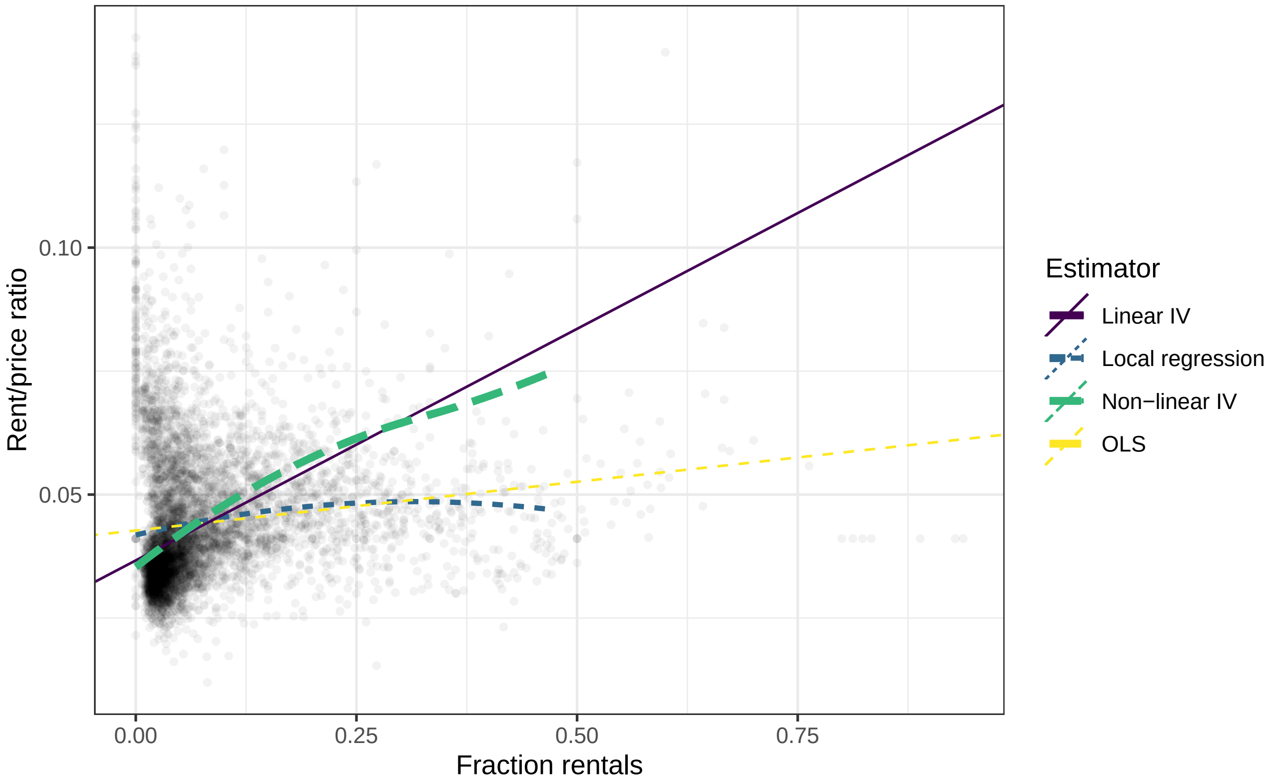

Note: The figure shows the estimated hurdle rate distribution (cdf). Every point corresponds to one housing type and shows its observed rent/price ratio and fraction of rentals. Our main estimator is the non-parametric non-linear IV (NPIV) by Chetverikov et al. (2018), where we instrument the fraction of rentals using the relationship between (Figure 4). The NPIV estimator excludes the smallest and greatest 1% of observations, which is why the NPIV fit is not available for the full support. In our counterfactuals, we use a constant extrapolation outside the estimated support. The other models are included for comparison. The linear IV model is a standard 2SLS regression, which imposes linear relations between the dependent variable, the independent variable and the instrument. OLS shows a linear regression and does not use any instrument. Local regression shows a B-spline regression and also does not use an instrument.

Figure D2: Counterfactuals with quality differences between rentals and owner-occupied houses

(a) Renting households

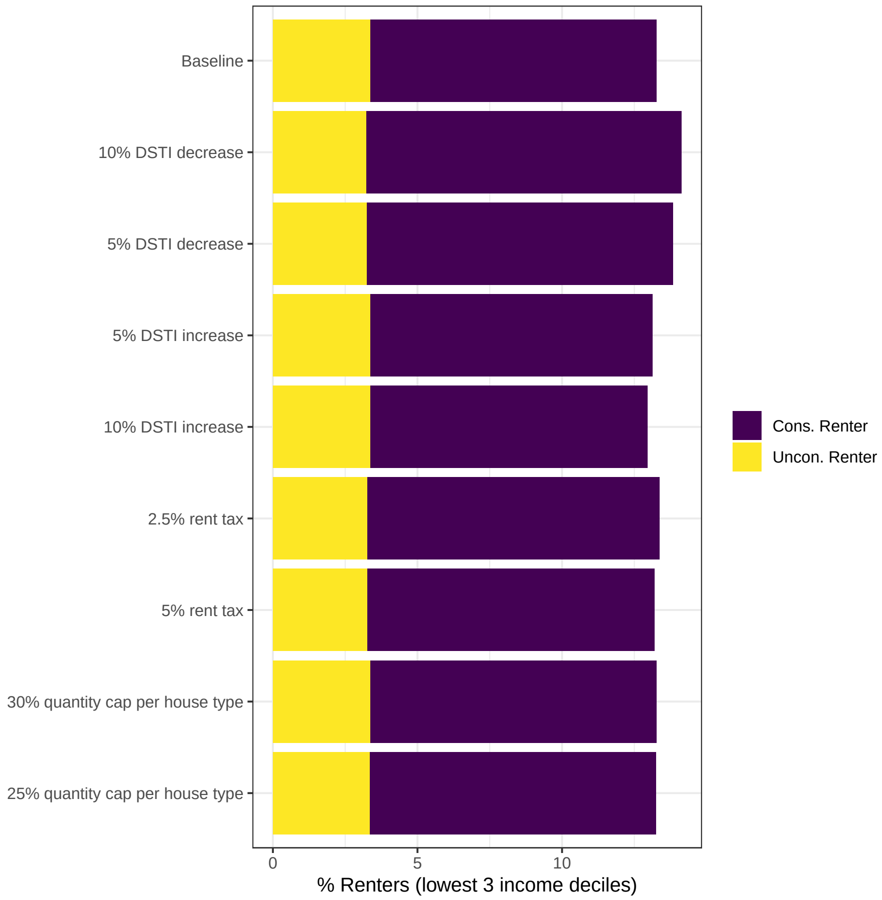

(b) Total surplus

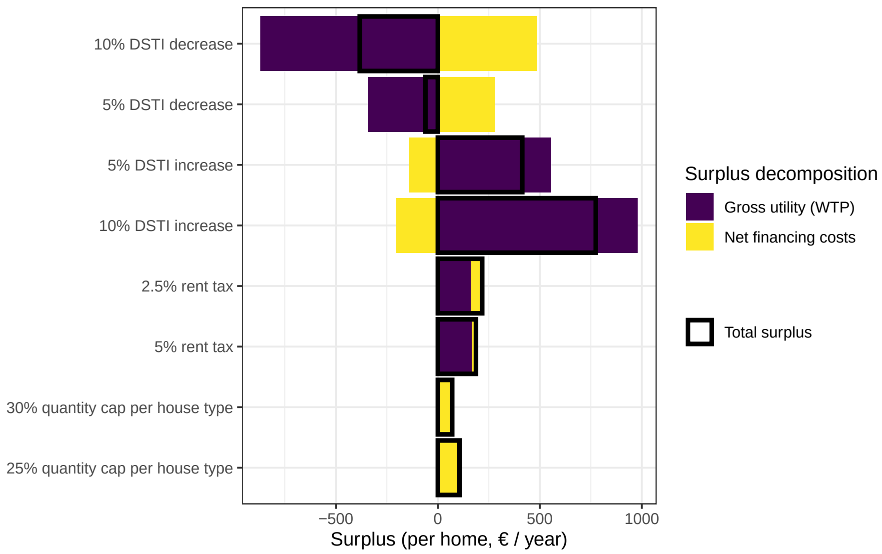

Note: The figures re-compute two main counterfactual outcomes using a demand function where we do not constrain rentals and owner-occupied houses to have the same average quality. The left panel is the equivalent of Figure 9, the right of Figure 10