The aggregate and distributional implications of credit shocks on housing and rental markets *

Juan Castellanos † Andrew Hannon ‡ Gonzalo Paz-Pardo §

April 2023

Abstract

We propose a joint model of the aggregate housing and rental markets in which both house prices and rents are determined endogenously. Households can choose their housing tenure status (renters, homeowners, or landlords) depending on their age, wealth, and income. We use our model to study the introduction in Ireland in 2015 of macroprudential policies that limited loan-to-value (LTV) and loan-to-income (LTI) ratios of newly originated mortgages. The introduction of stringent LTV and LTI ratios mitigates house price growth, but increases rents and reduces homeownership rates. As a result, middle-income households are negatively affected.

Keywords: House Prices, Rental Prices, Homeownership, Life-cycle, Housing Affordability, Credit Conditions, Macroprudential Policy

JEL classification: D15, E21, E30, E51, G51

*We would like to thank seminar participants at the ECB, the EUI, the St Gallen Macro Workshop, Norges Bank, BI Oslo and Paris School of Economics for their comments, and Jesus Bueren, Russell Cooper, Andrea Gazzani (discussant), Jonathan Heathcote, Ralph Luetticke, Ramon Marimon, Kurt Mitman, Ben Moll, Lukas Nord, Monika Piazzesi, Morten Ravn, Jirka Slacalek, Oreste Tristani, Gianluca Violante for helpful conversations. We are grateful to Matteo Crosignani for providing us with the distance data for our empirical analysis. The views expressed in this paper are those of the authors and do not necessarily reflect those of the European Central Bank.

†European University Institute, Department of Economics, Email: juan.castellanos@eui.eu.

‡University of Cambridge, Darwin College, Email: ah901@cam.ac.uk.

§European Central Bank, DG Research, Email: gonzalo.paz_pardo@ecb.europa.eu.

1. Introduction

Housing has a dual nature as an investment asset and as a consumption good. It represents the majority of wealth for most households and even those who don't own a home must consume housing services, which they usually obtain by buying them in the rental market. Because of their economic and political relevance, housing and rental markets are usually very closely regulated and subject to a wide range of tax advantages, subsidies and restrictions.

Understanding the dynamics of house prices and rental rates, together with their impact on household welfare, requires a joint study of both markets. Interventions that make renting less attractive might encourage households into homeownership, thus potentially raising house prices and lowering rental rates, whilst financial shocks that limit the availability of credit might have the opposite impact. The sizes of these effects depend on the segmentation between the markets for owner-occupied and buy-to-let properties (Greenwald and Guren, 2021), but also on the responsiveness of the landlord sector to rental and house prices and the differences in quality between owned and rented accommodation, amongst other factors.

In this paper, we develop an equilibrium model of the rental and housing sectors with several key features. First, households differ in their age, income, and wealth, and make endogenous housing tenure choices that lead them to be renters, homeowners, or landlords. Households can borrow through long-term mortgages for which downpayment and other constraints only bind at origination. Second, both rental and housing markets must be in equilibrium, which imply that house prices and rental prices must adjust to clear both markets as a result of potential shocks, but they may do so in different directions. This differs from standard assumptions in the macro-housing literature, in which the rental sector is either non-existent, has a fixed price, or is owned by a stylized foreign risk-neutral investor, thus implying that rents must either stay fixed or co-move with house prices. Third, we consider different scenarios for the convertibility of owner-occupied housing and rental accommodation, and show how they generate different implications for house prices, rents and their ratio.

We use our model to study the effects of a shock to the availability of credit for households, in the form of a borrower-based macroprudential policy intervention that imposes maximum loan-to-value (LTV) and loan-to-income (LTI) ratios to newly originated mortgages. We motivate our analysis with the reform introduced in Ireland in 2015, which set the minimum downpayment on a house to 20% and the maximum ratio of mortgages to household income to 3.5. Acharya, Bergant, Crosignani, Eisert, and McCann (2022) find that this reform reduced house price growth in areas where the limits were particularly binding. We extend their analysis to rents and find that it had the opposite effect: rental prices increased in those areas where the reform was more binding.

We show that we can rationalize this empirical observation through the lens of the model with the following intuition: in the presence of binding constraints to mortgage credit, prospective homeowners need to postpone or cancel their home buying decisions and stay renters for longer. Thus, more landlords need to enter the market, buy housing and provide it for rent. As long as the supply curve for rental housing at constant house prices is upward sloping, rents must go up to clear the market because the next marginal landlord must be compensated above the previous one in order to step in and provide additional rental housing. On the other hand, the effect on house prices depends on the relative qualities of rental and owner-occupied housing and the ease of convertibility amongst the two. If, as is true in the data for Ireland and most countries, rental housing is of smaller size or lower quality than owner-occupied housing, the increase in the share of renters will have a negative impact on average housing prices. In any case, the homeownership rate goes down.

These opposing effects on rental and house prices crucially hinge on the fact that we allow for *landlord heterogeneity*, a feature that is not present in models with a fully deep-pocketed and risk neutral investor who provides rental housing at user cost, nor in models in which the homeownership rate is fixed at 100% and there are no renters. The structure of our landlord sector is consistent with data from the Irish rental market, in which more than half of rental properties are owned by small landlords with one or two properties and less than 10 percent of properties are owned by large owners with 50+ properties (see Appendix E.1 for details).

Calibrating our model and replicating the Irish macroprudential reform, we find that it generates a 8.06% increase in rents and a 0.64% decrease in house prices. Moreover, it reduces the homeownership rate by 2.79 percentage points, with most of it being concentrated among households in the middle of the income distribution and between 35 to 55 years old. Finally, it also generates an increase in concentration in the housing and rental markets, with wealthier landlords holding a larger share of all rental properties (38% pre-reform vs. 39.3% post-reform). In terms of welfare, our model reveals that constraining housing credit is particularly harmful for renters, middle aged households and those in the middle of the income distribution because they find it more difficult to obtain a mortgage and therefore they need to postpone or cancel their buying decision while paying higher rental prices. Nonetheless, the aforementioned increase in rents and reduction of house prices benefit (prospective) landlords and households in the top of the income distribution since buying is cheaper, the constraints do not bind for these households, and the cash flow they obtain from their investment is higher.

Our results provide the first measure of the costs imposed on households by macroprudential policies through both the rental and housing sectors. However, given that we do not model aggregate uncertainty, cyclical buildup of risk or household's defaults, we cannot measure the benefits of those policies and, as a result, we are silent on optimal LTV or LTI ratios.

Although we have focused on the 2015 Irish macroprudential reform, our framework is more general and can be used to study the effects of other shocks in other economies. For instance, it suits itself naturally to study the effects on the housing and rental markets of monetary policy or fiscal policies such as mortgage subsidies, taxes associated with property purchases, etc. Landlord heterogeneity and the existence of many small landlords are not exclusive to the Irish economy, but a common feature of many European countries. In the United States, where institutional investors are quite developed, individual investors are still very relevant for the determination of housing and rental prices: on aggregate, private individual investors own 71.6% of all rental properties (Pew Research, 2021), with 14 million of them owning between 1 and 4 properties.

Related Literature Our paper contributes to studies of housing in macroeconomics (see Piazzesi and Schneider (2016) for a broad review), and in particular to recent papers that, inspired by the housing boom and bust associated to the 2008 global financial crisis, have analysed house price dynamics and their close relationship with credit, beliefs about future prices and fluctuations in liquidity. We contribute to it by focusing on the rental sector and studying its implications for welfare and aggregate dynamics.

In a general equilibrium model, Favilukis, Ludvigson, and Nieuwerburgh (2017) show that a relaxation of financing constraints can lead to a large boom in house prices. Justiniano, Primiceri, and Tambalotti (2019) focus on the role of the increase in the aggregate amount of credit supply in the economy, which they relate with the low mortgage rates that prevailed during this period. However, these models assume that all households are homeowners, and thus they lack the additional margin of adjustment through changes in housing tenure.

Kaplan, Mitman, and Violante (2020), instead, find that credit conditions had a very small role in house price movements. In their framework, the key driver of the movements in house prices and rents is the change in house price beliefs. They allow households to rent, but assume that a deep-pocketed rental sector is ready to provide rental housing at user cost, thus linking directly the evolution of rents and house prices. Their extension with household-landlords is closest to our framework, but they do not study its implications for rental rates and how they interact with house prices.

Garriga and Hedlund (2020) focus on the role of endogenous liquidity in housing markets and find that it amplifies the effect of credit shocks on housing prices. In their model, households can rent and the markets for rental and owner-occupied housing are segmented. In particular, they assume that very large dwellings are not available for rent, which in turn weakens the connection between rents and house prices present in the user cost framework.

In a recent paper, Greenwald and Guren (2021) rationalise the conflicting findings of these different contributions by showing that, in a model in which rental and owner-occupied properties are identical (no market segmentation), increases in homeownership rates do not impact house prices because households buy housing from deep-pocketed landlords that do not use credit. As a result, changes in credit conditions do not affect house prices. Instead, if the markets for rental and owner-occupied housing are segmented (i.e., they are different goods), the increase in demand by borrowers as a result of a credit shock raises house prices. Arslan, Guler, and Kuruscu (2020) follow this intuition and present a model with partially segmented housing markets and focus on the bank lending channel associated to credit supply shocks.

Our contribution builds on the intuition that assumptions about the rental sector are key for house price dynamics, but we extend it further by developing a model with full heterogeneity in which the decision to become a landlord is endogenous, as in Sommer and Sullivan (2018), who focus on the tax treatment of housing. Landlords in our model are heterogeneous in their age, income and wealth, which generates a positively sloping rental supply function at a given house price. Unlike Sommer and Sullivan (2018), our mortgages are long-term and have both loan-to-income and loan-to-value constraints that only bind at origination, which is key for the question we are considering. We find that, although the effects of a negative credit shock on house prices do depend on the assumptions about segmentation as shown in Greenwald and Guren (2021), the impact on rental rates is unequivocally positive.

Our modeling of housing and mortgages is based on Paz-Pardo (2021), which we extend to a general equilibrium setting where both house prices and rents are endogenously determined. We use this framework to connect the evolution of housing markets with distributional outcomes, such as Kiyotaki, Michaelides, and Nikolov (2011), focusing on an application related to macroprudential policy. The endogenous determination of rents and housing prices is novel to macro applications that study the effect of credit shocks, but is also present in state-of-the-art equilibrium models of local housing markets, such as Favilukis, Mabille, and Van Nieuwerburgh (2022), who focus on rent stabilization policies.

A broad literature has studied the benefits of macroprudential policies in terms of financial and macroeconomic stability from a modeling perspective, e.g. Lambertini, Mendicino, and Punzi (2013) or Farhi and Werning (2016). More recently, Ferrero, Harrison, and Nelson (2023) have also studied the joint optimal macro-prudential and monetary policy mix and put forward the idea of countercyclical borrower-based macro-prudential policies. Muñoz and Smets (2022) show that these countercyclical rules may trigger a reallocation of credit to the large real estate funds that operate in the rental market and consequently countercyclical LTV measures on institutional investors may be desirable for macroeconomic stability. Nonetheless, fewer have studied their negative consequences on household's welfare. One example is Queiró and Oliveira (2022) who study the effects of the loan-to-value and payment-to-income constraints implemented in Portugal in 2018 through a model à la Kaplan et al. (2020), and find that the reform is welfare reducing due to changes in homeownership and the quality of housing rather than through changes in prices.

Some other recent contributions have studied the effect of these reforms from an empirical perspective. Acharya et al. (2022) study the Irish reform in which we will base our empirical experiment, and using geographical heterogeneity in the intensity of the treatment find that the reform was successful in cooling the most tensioned housing markets, reallocating credit to rural areas and to higher-income borrowers. Van Bekkum, Irani, Gabarro, and Peydró (2019) study the introduction of a loan-to-value limit in the Netherlands and find that it reduced mortgage leverage, but also, consistently with the implications of our model, the share of renters that transition into homeownership.

Overview The rest of the paper is structured as follows. In section 2, we present the model. In section 3, we analyze the Irish macroprudential reform of 2015. First, by presenting some empirical evidence in section 3.1; and then by using a calibrated version of the model to analyze the effects on quantities and prices as well as on welfare. Parametrization is discussed in section 3.2, while model results are presented in section 3.3. Finally, section 4 concludes.

2. Model

In this section, we present a heterogenous-agent life-cycle model with housing and rental markets in which landlords choose their status endogenously. As a result, households in our model economy differ in their age, income, asset holdings and housing status. Renters and landlords meet in the rental market and determine the equilibrium rental rate. In the production side, there is a final good sector that produces non-durable goods and a construction sector that produces new houses.

2.1. Production

There are two production sectors in the economy: a final-good and a construction sector. They produce non-durable consumption goods and houses, respectively.

Final-good sector The competitive final-good sector operates a linear technology

$$Y_c = A_c N \tag{1}$$

where A_c is the constant aggregate labor productivity and N are the units of labor services. Profit maximization yields an equilibrium wage $w = A_c$.

Construction sector The competitive construction sector operates a Cobb-Douglas technology

$$Y_h = A_h L^{\alpha} S^{1-\alpha} \tag{2}$$

where L is the amount of buildable land (housing permits), S is the quantity of structures and $\alpha \in (0,1)$ is the constant share of land in production.

We assume that housing permits are fixed and are priced competitively, so that the housing developer solves the following static problem

$$\max_{S} \; p^h A_h \bar{L}^{\alpha} S^{1-\alpha} - p_L \bar{L} - S \tag{3}$$

where p_L is the equilibrium price of buildable land, and the price of structures (non-durables) is normalized to unity. Profit maximization implies the following housing investment function:

$$Y_h = A_h^{1/\alpha} \left((1-\alpha) \, p_h \right)^{(1-\alpha)/\alpha} \bar{L} \tag{4}$$

and consequently the elasticity of aggregate housing supply to house prices equals to $1/\alpha$. For a detailed derivation of equation (4) see Appendix A.2. As we shall see, equation (4) and the housing market clearing condition pins down the equilibrium per unit price of housing, p_h. However, this is not the final transaction price since we further assume that housing is discrete and it can be of two types: owner-occupied (OO) or buy-to-let (BLT) housing. As a result, we let the transaction price be a multiple of the per unit price: $p^h(\mathcal{H}_j) = \mathcal{H}_j p_h$ for $j \in \{oo, btl\}$ reflecting that these two types of housing are characterized by different qualities or sizes.

2.2. Households

Demographics Household's age is indexed by $j = 1, \ldots, J$. In the first $J^{ret} - 1$ periods they work. Thereafter they are retired until they die with certainty after age J.

Preferences Households derive utility from non-durable consumption and housing services. They value these streams of consumption according to

$$\mathbb{E}_0 \left\{ \sum_{j=1}^{J} \beta^{j-1} U\left(c_j, s_j(h)\right) \right\} \tag{5}$$

where $\beta > 0$ is the discount factor, $c > 0$ is consumption of non-durables and $s(h)$ is the consumption of housing services which depends on the type of the house where the household resides and acts as a multiplicative shifter. The per-period utility function is given by

$$U\left(c, s(h)\right) = \frac{\left(c\, s(h)\right)^{1-\gamma}}{1-\gamma} \tag{6}$$

where $\gamma > 0$ captures both risk aversion and intertemporal elasticity of substitution. As standard in the literature, we assume that the housing service flow for homeowners is larger than for renters. That is, $s(0) = 1$ and $s(h) > 1\ \forall h \in \{1, 2, \ldots, H\}$.

Endowments Working-age households receive an idiosyncratic labor income endowment. We assume that it has a deterministic component that depends on age and a stochastic, persistent component. That is

$$\log y = \log A_c + f(j) + \eta \tag{7}$$

where A_c is an index of aggregate productivity, $f(j)$ is a polynomial in age and η represents the stochastic persistent component of earnings. We estimate the earnings process non-linearly as in De Nardi, Fella, and Paz-Pardo (2020) (see Section 3.2.1 for details). Retired households receive a fixed fraction of their labor income in the last period prior to retirement for the rest of their lifetime.

Households are born with an initial endowment of liquid wealth that is drawn from a log-normal distribution. We also assume that they start their life as renters and thus have no housing wealth.

Liquid assets Households can save in a one-period risk-free bond, $a \geq 0$ that yields a constant interest rate, r, which is determined in the world market and therefore is exogenous.

Housing choices Households can also save through buying houses, which serve a dual role as assets and source of consumption flows. Households who do not own a house must rent one in the market at the unit rental rate p^r. Homeowners do not have to pay rent. When a homeowner buys additional houses as an investment, she rents them out in the market and receives p^r per period per house. Note that these assumptions allow us to distinguish between renters ($h = 0$), homeowners ($h = 1$) and landlords ($h > 1$). As anticipated above, we further assume that there are two types of houses which differ in their quality $\mathcal{H}_j$ and that the final transaction price payed or received by households depends on it, i.e. $p^h(\mathcal{H}_j) = \mathcal{H}_j p^h$.

Housing is illiquid since buying and selling housing units is costly. Thus, we assume that households pay a proportional transaction cost, $\tau^h p^h(\mathcal{H}_j)$, that depends on the value of the house being bought or sold. These capture real estate agent fees, taxation and other administrative costs. Houses are also costly to maintain. Therefore, home-owners and landlords pay maintenance costs to keep up with their depreciation, $\delta^h p^h(\mathcal{H}_j)$, where δ^h is the depreciation rate.

Mortgages The purchase of a house can be financed through a mortgage at a fixed rate $r(1+\kappa)$, where $(1+\kappa)$ is the is the intermediation wedge between the mortgage rate and the risk-free rate. To reduce the dimensionality of the problem, we treat mortgages as negative asset holdings $a \leq 0$, which prevents mortgagors from simultaneously having liquid assets. In other words, a denotes the net asset position.

At origination, the borrower must satisfy two constraints. First, a maximum loan-to-value (LTV) limit, which imposes that the size of the mortgage has to be smaller than a fraction of the value of the house. And second, a loan-to-income (LTI) requirement that requires that the household does not borrow more than a multiple of its current income. Formally,

$$a' \geq -\lambda_{LTV}\, p^h(\mathcal{H}_j)\, h' \tag{8}$$

$$a' \geq -\lambda_{LTI}\, y \tag{9}$$

where $\lambda_{LTV} < 1$ and $\lambda_{LTI} > 1$. As in reality, these two constraints do not bind after origination, so households do not need to renegotiate their mortgages every period and thus are not subject to these constraints every period.

After origination the borrower chooses the repayment schedule freely. Nevertheless, we impose that all debts must be paid before the terminal age J, i.e. $a_J = 0$ and that a minimum amortization payment must be made in each period, together with interest payments. As in Kaplan et al. (2020), the minimum payment is determined by the constant-amortization formula

$$m_j = \frac{r\,(1+\kappa)\,(1+r\,(1+\kappa))^{J-j}}{(1+r\,(1+\kappa))^{J-j} - 1} \tag{10}$$

Household's Problem A household of age j, income y, with h houses and a assets solves the following dynamic program

$$V(a,h,y,j) = \max_{a',h'} \left\{ \frac{(c\ s(h))^{1-\gamma}}{1-\gamma} + \sigma_\varepsilon \varepsilon(h) + \beta \mathbb{E} V\left(a', h', y', j+1\right) \right\}$$

$$\begin{aligned}
&\text{s.t.} \\
&c + a' + p^h(\mathcal{H})h' + \tau^h p^h(\mathcal{H})|h'-h| + \delta^h p^h(\mathcal{H})h \leq \\
&\quad y + \left(1 + r\left(1 + \mathbb{I}_{a'<0}\kappa\right)\right) a + p^h(\mathcal{H})h + p_r(h-1) \\
&a' \geq \begin{cases} \max\left\{-\lambda_{LTV}\ p^h\left(\mathcal{H}\right)\ h', -\lambda_{LTI}\ y\right\} & \text{if } h' > h \\ a(1 + r(1+\kappa) - m(j)) & \text{if } h > 0 \text{ and } a < 0 \\ 0 & \text{otherwise} \end{cases}
\end{aligned} \tag{11}$$

where $\sigma_\varepsilon \varepsilon(h)$ are choice-specific random taste shocks that are *i.i.d.* Extreme Value Type I distributed with scale parameter σ_ε. These represent shocks to the utility of homeownership (i.e., they are alike to moving shocks), but are also computationally convenient as they help to smooth out expected value functions (Iskhakov, Jørgensen, Rust, and Schjerning, 2017).

2.3. Equilibrium

A competitive equilibrium in this economy consists of a value function $\{V\}$, policy functions for the households $\{c, a', h'\}$, policy functions for the firms $\{N, L, S\}$ and prices $\{w, p_h, p_r, p_l\}$ that jointly solve the household, final-good firm and construction firm problems, and satisfy the following market clearing conditions

$$\text{Housing production:} \quad Y_h = (\delta_h + 1/T)H \tag{12}$$

$$\text{Aggregate housing:} \quad H = \sum_{n=1}^{N} \mathcal{H}_n H_n^{sh} \quad \text{where} \quad \sum_{n=1}^{N} H_n^{sh} = 1 \tag{13}$$

$$\text{Total production:} \quad Y = Y_c + Y_h \tag{14}$$

$$\text{Resources:} \quad Y = C + S \tag{15}$$

for a given risk free rate r.

3. A macro-prudential reform: the case of Ireland

Macro-prudential regulations that limit household leverage in the residential mortgage market have been widely used by policymakers to smooth the house price and credit cycles. We study the case of Ireland, whose central bank introduced these measures for the first time in February 2015 after a first discussion in October 2014. In particular, the Central Bank of Ireland established a maximum LTI limit of 3.5 and several LTV limits depending on the borrower type. For primary dwellings (PD) the limit was set to 80% of the value of the house; for first-time-buyers (FTBs), the limit was more generous: 90% for the first 220,000 € and 80% for the excess amount; and for buy-to-let (BTL) properties the threshold was more stringent and set to 70%. Additionally, a share of the total amount of lending offered by banks was exempt of the limits (15% of all lending for LTV limits and 20% of all lending for LTI limits).

The prompt implementation of the reform, paired with data availability, makes Ireland a compelling case study. We now turn to showing the effects that the reform had on house prices and rents in the data, which we will use as motivating evidence for the calibrated version of our model.

3.1. Empirical evidence

Using data for the universe of originated mortgages in Ireland, Acharya et al. (2022) study the 2015 reform and find that it generated a reduction in house price growth. In order to control for potentially confounding effects in macroeconomic aggregates, they develop a *distance* measure that correlates with exposure to the macro-prudential reform. In counties where house prices were high with respect to incomes, many mortgages were originated before the reform at or above the limits: these are *low-distance* areas where the reform is expected to have stronger effects. Instead, in counties where house prices were relatively low with respect to incomes, the reform was closer to non-binding and thus expected to have low to no effects. Consistently, they find that the distance measure positively correlates with house price growth around the reform: house prices grew more in areas where the constraints were less binding, while house price growth moderated in areas where the intervention was more binding.

We extend their framework to look at the effects on rents. We use the distance measure they computed and merge it with house price and rental data extracted from daft.ie (Lyons, 2018). Following their empirical implementation, we measure changes in house prices and rents between the third quarter of 2014 and the last quarter of 2016 and regress them on the aforementioned distance measure. That is,

$$\Delta HP_i = \beta_0 + \beta_1 \text{Distance}_i + \epsilon_i \tag{16}$$

$$\Delta HR_i = \gamma_0 + \gamma_1 \text{Distance}_i + \nu_i \tag{17}$$

Table 1 shows the results. We replicate the positive coefficient they obtain for house prices, but we find that the effect on rents was the opposite: rents increased by more in areas where the macro-prudential intervention had stronger effects. This movement of rents and house prices in opposite directions is restricted to the time around the implementation of this policy and is not present for the rest of our sample period. Figure 1 shows the ratio of coefficients β_1 and γ_1 in Equations (16) and (17) if we estimate them in different sample periods (x-axis). Specifically, we keep the same time window as in our main regression (9-quarter changes) and report the coefficients as a function of the central part of the window. We keep the distance measure fixed to its value in Acharya et al. (2022), which corresponds to 2014.

Table 1: Effect of lending limits on house and rental prices

	Δ House prices	Δ Rents
Distance	0.289	-0.171
	(0.068)	(0.039)
Obs.	54	54
R^2	0.34	0.31

The unconditional average and median of this coefficient is positive, reflecting that, in general, house prices and rents tend to comove, which is consistent with theories that set the value of a house to the expected discounted value of future rents. However, around our credit shock the fact that there are constrained renters who can't buy breaks this long-run relationship and generates effects going in opposite directions. The fact that this negative coefficient is limited to the time around the reform is a model implication, and is reassuring that the distance measure does not capture other relevant, time-invariant omitted variables (such as urban vs. rural, for example) which were themselves driving the relationship between house prices and rents.

In short, this placebo test confirms that our findings are not driven by time-invariant omitted variables which are correlated with the distance measure and reinforces the idea that the credit shock induced a decoupling of the usually positive relationship between the evolution of house prices and rents.

Figure 1: Ratio of coefficients for house prices and rents, by year

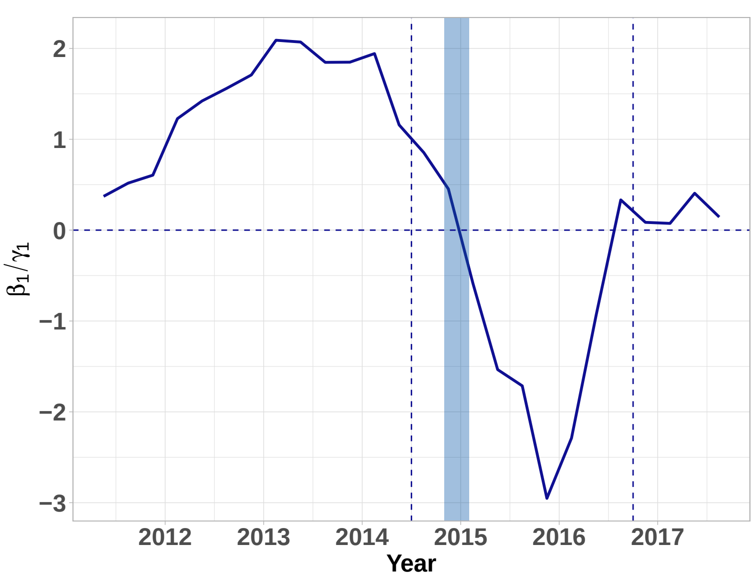

3.2. Parametrizing the model

In this section, we parametrize our model to be consistent with the cross sectional features of the Irish economy and use it for a twofold purpose. First, to understand the opposite effects on house prices and rents of the macroprudential reform; and second, to further analyze the distributional effects and the costs imposed on households by these reforms, taking into account a broad life-cycle perspective.

As standard in the macroeconomic literature, we assign some of these parameters externally, while others are chosen internally with the objective of minimizing the distance between a collection of data and model moments.

3.2.1. Earnings Process

We estimate our earnings process on a combination of Household Finance and Consumption Survey (HFCS) data and European Union Statistics on Income and Living Conditions (EU-SILC) data. The HFCS, which takes place every three years, collects rich data on the income and wealth of European households, including their homeownership status, rental income, etc., which we also use as targets for our model. We use the HFCS to extract the average age-earnings profile in the Irish economy after taking into account year effects. However, in the case of Ireland the HFCS lacks a panel component, so we resort to EU-SILC data to study the annual fluctuations of income. Although it is targeted to produce statistics on poverty and living conditions, it is nationally representative. We use the longitudinal version of the survey (between 2005-2019) and after dropping households with low attachment to the labour force (below 1,500 € per year of labour market income) we are left with 1,270 observations to estimate our earnings process.

Our measure of income in the data is disposable household income after both taxes and transfers. To filter out potential measurement error in the survey whilst keeping potentially relevant rich features of earnings dynamics, we use the procedure described in Arellano, Blundell, and Bonhomme (2017) to extract out a transitory component and keep the persistent component of earnings.

Applying this procedure allows us to estimate earnings dynamics under flexible assumptions, and in particular incorporating potential age-dependence, non-normalities and non-linearities in earnings dynamics. The first element is of particular relevance for our question. Most households become homeowners when they are relatively young, still changing jobs and potentially subject to large fluctuations to their labour market income. A standard earnings process in which earnings are a random walk is a poor representation of the earnings risk faced by households at this particular age. Middle-aged households with stable jobs, instead, have much higher persistence, but significant negative skewness risk (e.g., through job loss). For a detailed description of the method and the economic implications of flexible earnings dynamics, see De Nardi et al. (2020).

3.2.2. Externally calibrated parameters

The model period is one year. Households enter the economy at age 25, they retire with certainty at age 65 and live until age 95. This means that $J^{ret} = 41$ and $J = 71$. There is no population growth. We set the CRRA risk aversion coefficient, γ, to 2, a common value in the literature. The scale parameter of the taste shock, σ_ε, is within the range suggested by Iskhakov et al. (2017) and equal to 0.05.

The house qualities, $\{\tilde{h}_{oo}, \tilde{h}_{btl}\}$, are chosen to replicate the owner-occupied to buy-to-let house price ratio, which is 1.21 in the data. The housing depreciation rate, a cost borne by households, is set to be 1.2% per year and it is within the range of typical values used in the literature. The transaction cost for selling or buying a house τ^h equals 4% of its value. The maximum loan-to-value, λ_{LTV}, and loan-to-income, λ_{LTI}, ratios are 1.0 and 6.0, respectively. This is consistent with the evidence in Kelly, McCann, and O'Toole (2018) that estimate the $98th$ percentile of observed LTI and LTVs ratios on quarterly mortgage data during the period 2003 to 2011.[1]

The risk-free rate on liquid savings r is set to 2%. The mortgage rate equals the risk-free rate times an intermediation wedge $(1+\kappa)$, which takes a value of 1.02, i.e. we assume a 2% spread on mortgage borrowing.

The final good aggregate productivity shifter, A_c, is set to 1.25. The share of land, α, used in production in the housing sector is fixed to 0.33 and the amount of buildable land, $\bar{L}$, is equal to 1.

Table 2: Parameter values

Parameter	Interpretation	Value
Externally calibrated:		
J^{ret}	Working life (years)	41
J	Length of life (years)	71
γ	Risk aversion coefficient	2.0
σ_ε	Scale parameter (taste shock)	0.05
$\{\tilde{h}_{oo}, \tilde{h}_{btl}\}$	Housing qualities	$\{1.036, 0.8562\}$
δ^h	Housing depreciation rate	0.012
τ^h	Proportional transaction cost	0.04
λ_{LTV}	Maximum loan-to-value ratio	1.0
λ_{LTI}	Maximum loan-to-income ratio	6.0
r	Risk-free rate	0.02
κ	Intermediation wedge	0.02
A_c	Aggregate labor productivity	1.25
$\bar{L}$	Amount of land	1.0
α	Share of land in production	0.33
Internally calibrated:		
β	Discount factor	0.9375
$s(h>0)$	Utility premium from ownership	1.6
A_h	Scaling factor in housing production	0.12

[1]Note that prior to the 2015 reform, there was no institutional limits. Therefore, these limits were unofficially imposed by Irish banks and were based on their own risk assessment.

3.2.3. Internally calibrated parameters, targets and model fit

The remaining three parameters: the discount factor β, the homeownership utility premium $s(h > 0)$, and the scaling factor in housing production A_h, are jointly set to match four moments of the data. In particular, we try to match the average wealth to income ratio, which is around 7 in the HFCS; the homeownership rate that was around 80% before the 2008 crisis according to EU-SILC; a house price to income ratio of 5 that is consistent with the data in the Central Statistics Office (CSO); and the house price to rent ratio that is computed using data from the Residence Tenancies Board (RTB) and the CSO.

Table 3: Targeted and untargeted moments

Moment	Model	Data	Source
Targeted:			
Wealth to income ratio	5.32	6.78	HFCS
Homeownership rate	79.13%	80%	EU-SILC
Avg. house price to income ratio	4.90	5.0	CSO
House price to rents ratio	23.00	22.58	RTB/CSO
Untargeted:			
Rents to avg. income ratio	0.2132	0.2216	RTB/CSO
Share of households with 3+ properties	3.97%	5.11%	HFCS

The first block of Table 3 shows the exact value of these moments in the data as well as their model counterparts which are obtained using the parameters reproduced in the last block of Table 2. The model is able to match the average homeownership rate, the average house price to income ratio, and the house price to rent ratio reasonably well. However, it under-predicts the average wealth to income ratio since we do not model the stock market nor include bequest motives.

The model is also able to replicate the share of landlords in the economy, both at the aggregate level and along the age distribution. In the model, the share of landlords with two rented out properties (the upper bound) is 3.97%, while in the data 5.11% of landlords own two or more rental properties. The model also does a great job in capturing the life cycle dynamics as shown in Figure 2. It is only at the middle of the life-cycle when some households can afford a second or third home, which later in old-age they sell to finance retirement. In both the model and the data, the share of renters decreases with age, and many households are able to access homeownership earlier in life. During the retirement period, the model overpredicts the share of renters and the extent to which landlords sell their properties to finance retirement. This mismatch is a standard feature of life-cycle models who do not model a set of relevant features of retiree saving behaviour, including precautionary savings related to medical costs or long-term care, bequest motives, etc. (Nakajima and Telyukova, 2020).

In any case, these life-cycle patterns are endogenously captured by the model without explicitly targeting them, which is reassuring about the validity of the model as a laboratory to study the distributional effects of the macro-prudential reform discussed above.

Figure 2: Housing status (number of properties)

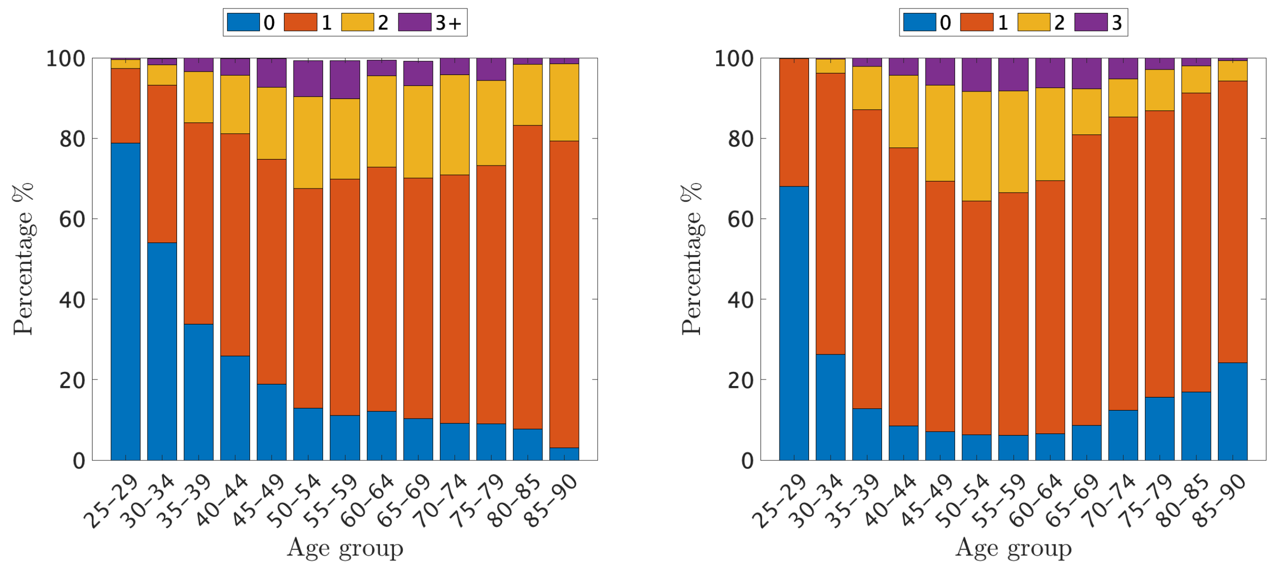

(a) Data (b) Model

3.3. The reform through the lens of the model: exploring the mechanisms

3.3.1. Intuition: a supply and demand explanation

Before we turn into the results of the model, i.e. how the Irish macroprudential reform affected house and rental prices, homeownership rates and welfare, we present the intuition through a supply and demand illustration. We use the share of renters (demand) or the share of landlords adjusted by the number of properties they own (supply) on the x-axis and the rent-to-price ratio on the y-axis.

As shown in Figure 3, rental demand is downward sloping because increases in the rental to house price ratio incentivize homeownership and consequently less and less households are willing to be renters. On the other hand, such increase makes buying buy-to-let properties more attractive and more and more households are willing to be landlords. This results in an upward sloping rental supply curve. As standard, the intersection of these two curves form an equilibrium which determines the price. As illustrated in panel (a), such equilibrium requires that the share of renters equals that of houses offered by landlords in our model economy, which in turn determines the rental rate and the average house price.

Now, consider the impact of a credit reduction associated for example to the introduction of macroprudential mortgage limits. On impact the reform primarily affects potential buyers that were close to the borrowing limit before and that now do not qualify for a mortgage. These households are not able to buy a house anymore, and therefore need to rent. This shifts the demand curve outwards as shown by the blue dash line in panel (b) of Figure 3.

Figure 3: Supply and demand in the rental market

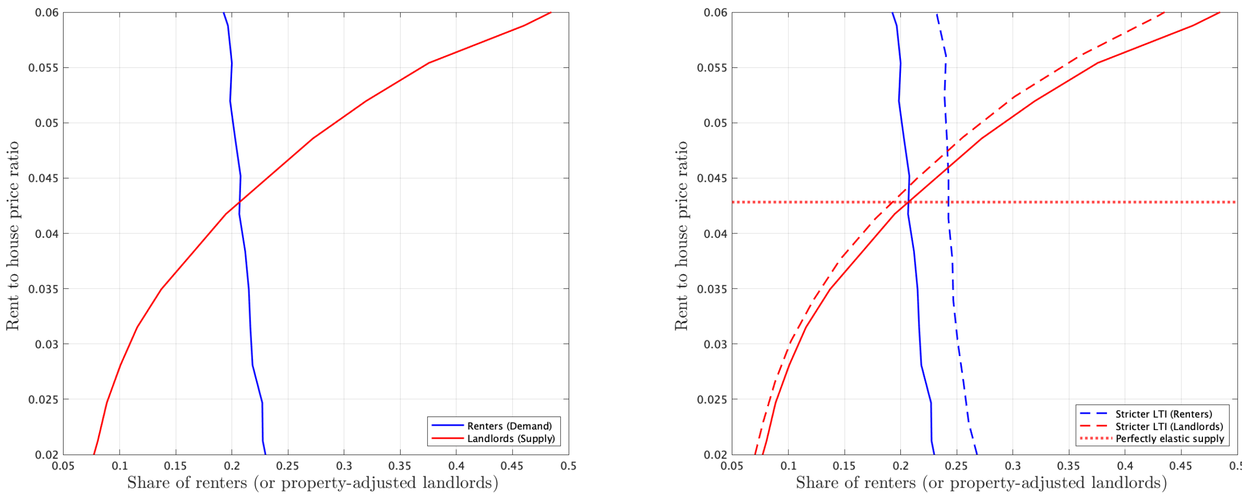

(a) Equilibrium

(b) Credit shock

In a model with perfectly elastic rental supply (red dotted line), the increase in rental demand only translates into a reduction of the homeownership rate since the share of renters goes up. Prices do not move because deep-pocketed landlords are willing to buy as many houses as needed at the present value of rents to meet rental demand. As shown in Greenwald and Guren (2021), this is why in the baseline model of Kaplan et al. (2020) credit conditions do not affect house prices.

As discussed above rental supply is upward sloping in our model. Consequently, an increase in rental demand associated with the reduction in credit results not only in changes in the homeownership rate but also in an increase in the rent to price ratio. This increase in the price ratio is slightly amplified in our model because landlords also use credit to buy additional rental properties, which slightly shifts rental supply inwards (red dashed line).

Overall, a reduction in credit results in an increase of the rent to price ratio and a reduction of the homeownership rate because rental demand shifts outwards, rental supply is upward sloping and (potential) landlords are a bit sensitive to credit conditions.

3.3.2. Steady state comparison

We first study the aggregate effects of the macroprudential reform if it were to be permanent. We do so by comparing two steady state equilibria that only differ in their institutional parameters. First, the LTI and LTV limits are those specified in Section 3.2 for the *pre-reform* economy. Second, the model is solved using the LTI and LTV limits established by the Central Bank of Ireland in 2015 for the *post-reform* economy. In particular, we simplify some of its aspects and directly impose $\lambda_{LTI}^{post} = 3.5$ and $\lambda_{LTV}^{post} = 0.8$. Results are shown in the first two columns of Table 4.

Table 4: Quantity and price effects across steady states

	Pre-Reform	Post-Reform	Only LTI
Rent-to-Price	4.38%	4.73%	4.73%
Average house price to income	4.90	4.87	4.87
Rent to Income	0.21	0.23	0.23
Homeownership rate	79.13%	76.34%	76.36%
Share of households with 3 properties	3.97%	4.65%	4.66%
Share of houses in hands of 3-property landlords	38.06%	39.29%	39.47%

The more stringent borrowing limits lead to an increase in the rent-to-price ratio and to a reduction of the home-ownership rate in the long run, as shown graphically above. In particular, the homeownership falls 2.79 percentage points while the rent-to-price ratio increases by 8.76%. As we showed empirically in Section 3.1, the macroprudential reform had opposite effects on house prices and rents. Consequently, the increase in the rent to price ratio could be originated by an increase in rents or a drop in house prices. Importantly, our model is able to disentangle the two. We find that rental rates rise by 8.06% while house prices fall by 0.65%. Thus, our model is consistent with the negative empirical correlation of macroprudential borrowing limits and house prices, as well as the positive empirical correlation of these limits and rental prices.

Figure 4: Homeownership rate

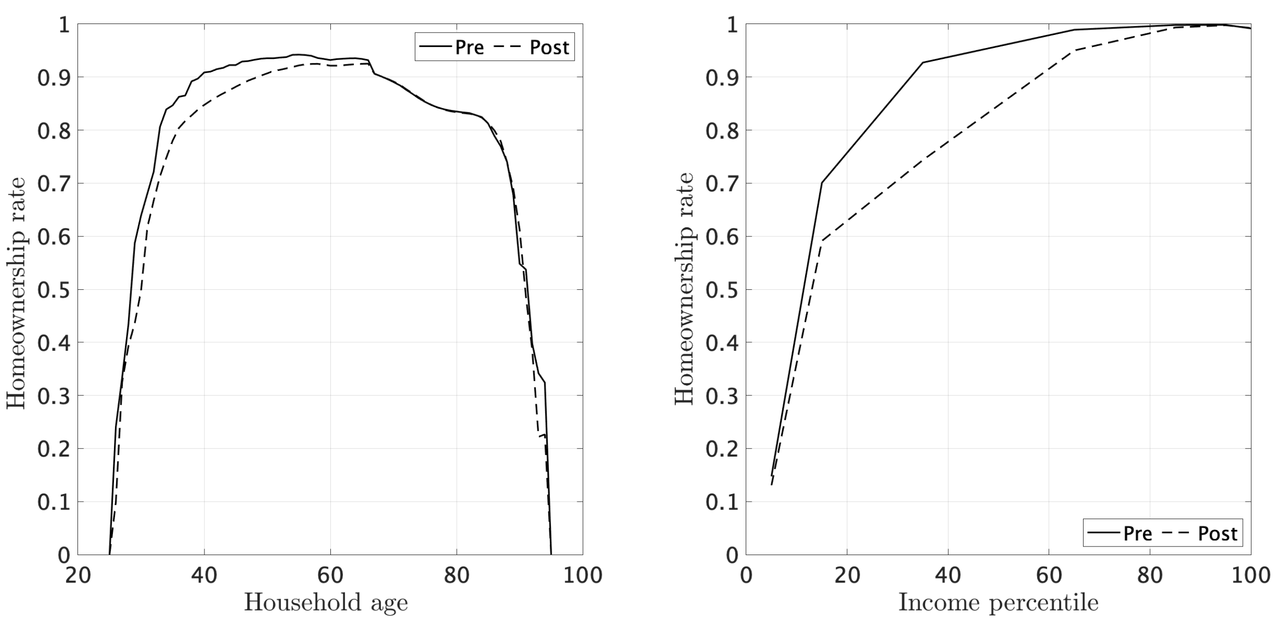

(a) Along the age dist.

(b) Along income dist.

Regarding the fall in the homeownership rate, we find that it is concentrated among middle aged households and those in the middle of the income distribution (see Figure 4). These results are very intuitive. Very poor households cannot afford to purchase a house independently of the credit conditions while the very rich are typically unconstrained or they do not need a mortgage. As a result, the reform has very little to no impact on their choices. However, for middle income households, that typically finance their housing purchases through mortgage borrowing, the tightening of these conditions have direct/partial equilibrium effects because they do not longer qualify for a mortgage and indirect/general equilibrium effects because higher equilibrium rental prices make it harder to save for downpayment.

Since there is a reduction in the homeownership rate, the share of renters must rise in the *post-reform* economy. As shown in the last two rows of Table 4, this extra demand for rental units is met by existing and new landlords. In fact, the share of landlords with two rental properties increases by 0.68 p.p., and the share of rental properties in hands of the wealthiest landlords rise from 38.06% to 39.29%. Thus, the reform lead to a higher degree of concentration in the rental market.

Finally, we decompose the effects of each of the two borrowing limits by solving for a third steady state equilibrium in which only the new LTI limit is imposed. Results are shown in the third column of Table 4. Given the similarity of this equilibrium with the *post-reform*, we can conclude that most of the effects are originated by the tighter LTI limits since it is the one that typically binds given the Irish institutional framework and parametrization.

3.3.3. Welfare effects along the transition

We now consider the effects of the transition from the (initial) *pre-reform* steady state to the (final) *post-reform* steady state. That is, agents unexpectedly observe that LTV and LTI limits become more stringent with the LTI limit dropping from 6 to 3.5 and the LTV limit falling from 100% to 80%. The reduction of these limits is permanent and it is perceived as such by the agents in the economy.

Transition paths Figure 5 show the transition paths for rents, average house prices and homeownership rates in deviations from the initial steady state. After the introduction of the regulatory LTV and LTI limits, homeownership rate falls by 1.8 p.p. upon impact since there is a significant mass of households that at the new limits do not qualify for a mortgage and consequently are forced to obtain housing services in the rental market. This increase in rental demand puts upward pressure in rental prices which immediately rise by 6%. Similarly, average house prices also react upon impact and fall by 1.7%. These price changes incentivize (prospective) landlords to buy additional rental houses, bringing back up the average house price to its new steady state level, which is 0.6% lower than in absence of the reform. In any case, rental prices continue adjusting upwards in subsequent periods because saving for downpayment is more difficult as a bigger fraction of income is destined to obtain housing services in the rental market, which in turn forces homeownership to keep adjusting downwards putting upward pressure on rents despite the slow increase in supply, specially coming from new landlords.

Figure 5: Transition dynamics

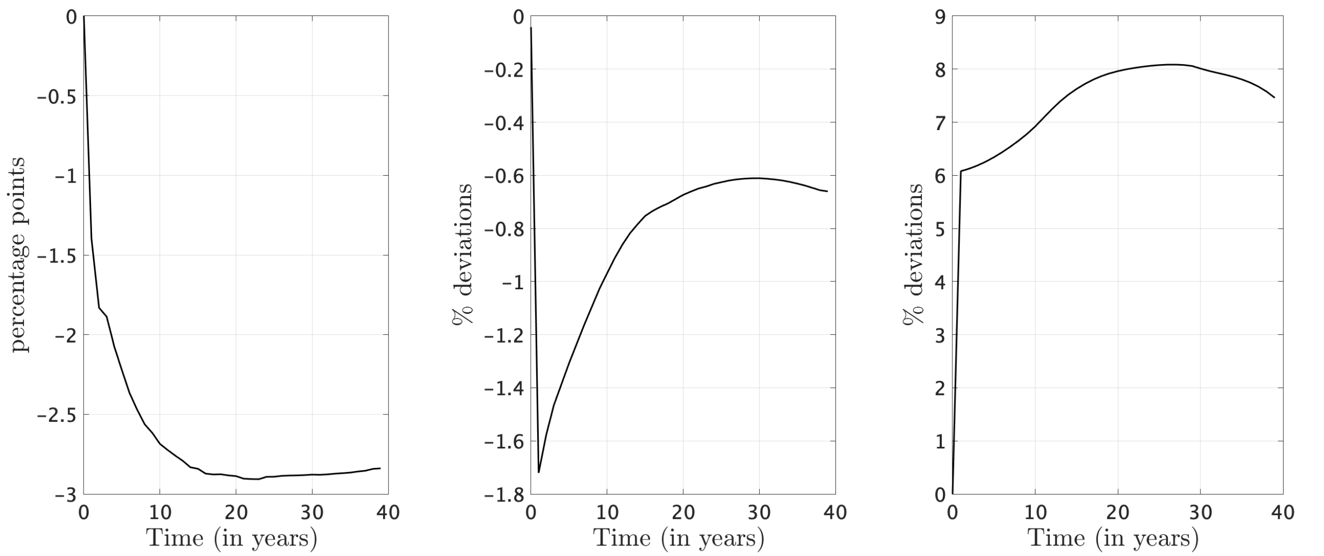

(a) Homeownership rate (b) Average house price (c) Rents

In Figure 10 in the appendix we also depict the evolution of the share of landlords with 2 or 3 properties as well as that of owner-occupied and buy-to-let housing prices.

Welfare We evaluate the distributional effects of the macroprudential reform through the traditional lifetime consumption equivalent variation (CEV) measure. This metric informs us about how much consumption (in percentage) needs to change in the pre-reform economy such that the households are indifferent between living in the pre-reform steady state and living through the transition induced by the policy reform. Formally, for a given set of state variables $x = (a, y, h, j)$, the consumption equivalent variation $g(x)$ is computed as

$$V_0(x; g) \equiv (1 + g)^{1-\gamma} V_0(x) = V_1(x) \quad \implies \quad g(x) = \left[\frac{V_1(x)}{V_0(x)}\right]^{\frac{1}{1-\gamma}} - 1 \tag{18}$$

where we are using the fact that the utility function is CRRA. From (18) it is easy to realize that a negative value of $g(x)$ is associated with agents being worse-off by the introduction of the reform.

Figure 6 shows this metric along the age distribution (left) and along the income distribution at the initial age (right) for three different scenarios. In the first scenario, we compute the CEV at constant prices. That is, we fix both rental and house prices to its pre-reform values and evaluate the welfare effects of constraining mortgage borrowing along the transition. As expected, such intervention is welfare reducing. More interestingly, and as shown by the dashed line, these losses are particularly large for the young and those in the middle of the income distribution.

Figure 6: A decomposition of the lifetime CEV – the role of rental prices

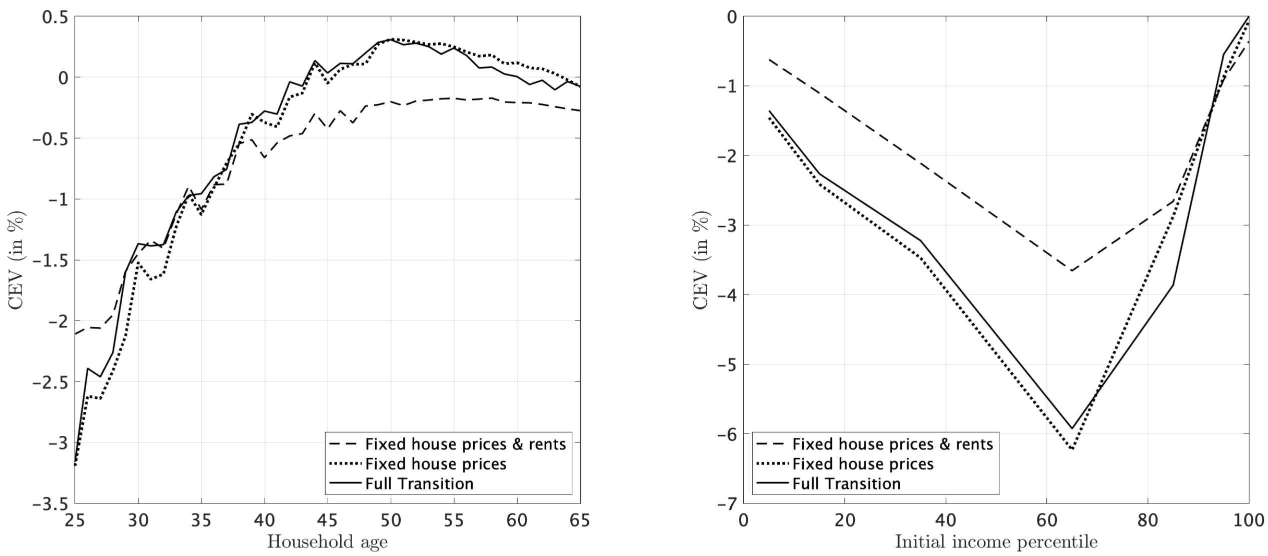

(a) Age

(b) Income

The second scenario, depicted by the dotted line, shows the effects of rental price movements. In particular, to evaluate its welfare effects, on top of those induced by changing the LTI and LTV limits, we feed in the equilibrium path for rents while still keep the house prices at their pre-reform values. As revealed by the distance between the dashed and dotted lines, the sharp jump in rental prices harms the young who are typically renters, while it benefits the middle-aged who are more likely to be landlords. Along the income distribution, the change in rental prices severely affects those in the bottom and middle of the income distribution since they must pay higher rental prices. Moreover, prospective homeowners will need to cancel or postpone their house buying decisions since rental payments constitute now a larger fraction of their expenses, making impossible, in some cases, to save enough for downpayment.

Finally, the third scenario shows the overall welfare effects in which both rental and house prices behave as in Figure 5. This is depicted by the solid line, which follows closely the dotted line. That is, the drop in house prices in isolation does not have a big impact on average welfare by age. Nevertheless, it has more clear distributional implications along the income distribution: those in the middle-lower part of the income distribution are better off because they can now buy a house at a cheaper price, while those at 70-90 percentiles of the income distribution are slightly worse off since their houses have lost some value, and therefore provide less cushion for unexpected income shocks.

The aforementioned decomposition of welfare effect has highlighted the importance of housing status on the bearing of the welfare costs of the macro prudential reform. Therefore, we also analyze the welfare effects for renters, homeowners and landlords in Figure 7. This figure confirms that the macroprudential reform benefits landlords, harms renters and has a small impact on homeowners on average.

Figure 7: CEV – the role of housing tenure

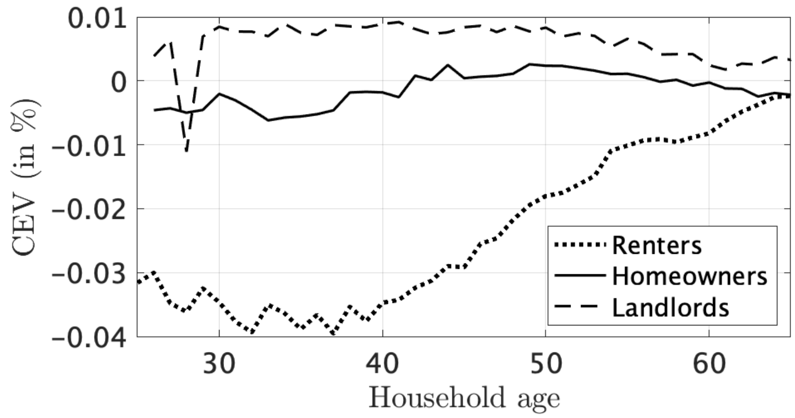

4. Conclusion

In this paper, we empirically show that the introduction of LTI and LTV limits in Ireland in 2015 had opposite effects on house prices and rents. To rationalize this finding, we build a life-cycle incomplete markets model with *heterogenous landlords*, which allows us to disentangle the effects of these borrowing-based macroprudential policies on both housing and rental markets.

Viewed through the lens of the model, the 2015 Irish macroprudential reform led to the intended reduction in house prices, but at the cost of increasing rental rates and reducing the homeownership rate. Moreover, it also increased the concentration in the rental market as the share of rental housing in hands of the wealthiest landlords increased.

These aggregate effects impacted households differently. The young, the middle-income and the renters suffer the most from the reform because they have to pay higher rents and are forced to postpone or cancel their buying decisions. On the other hand, top income earners, soon-to-be retirees and landlords slightly benefited because these agents are not constrained by the new borrowing limits and can buy additional housing at lower prices.

Our model does not include aggregate uncertainty, cyclical buildup of risk or household default. Embedding our framework in a model that incorporates the benefits of macroprudential regulations would be an interesting next step with which to study optimal LTV and LTI ratios. Besides, in our model all of the heterogeneous landlords are relatively small. As Muñoz and Smets (2022) show, the role of institutional investors matters to study the effects of macroprudential policies. Thus, a fruitful area for future work would be to study the relative role of these two different types of investors, their heterogeneity across housing markets, countries and over time, and their macroeconomic implications.

References

Acharya, V. V., K. Bergant, M. Crosignani, T. Eisert, and F. McCann (2022): "The Anatomy of the Transmission of Macroprudential Policies," *The Journal of Finance*, 77, 2533–2575.

Arellano, M., R. Blundell, and S. Bonhomme (2017): "Earnings and Consumption Dynamics: A Nonlinear Panel Data Framework," *Econometrica*, 85, 693–734.

Arslan, Y., B. Guler, and B. Kuruscu (2020): "Credit supply driven boom-bust cycles," *BIS Working Papers*, No 885.

De Nardi, M., G. Fella, and G. Paz-Pardo (2020): "Nonlinear Household Earnings Dynamics, Self-Insurance, and Welfare," *Journal of the European Economic Association*, 18, 890–926.

Department of Finance (2019): "Institutional Investment in the Housing Market," Irish Government Report.

Farhi, E. and I. Werning (2016): "A theory of macroprudential policies in the presence of nominal rigidities," *Econometrica*, 84, 1645–1704.

Favilukis, J., S. C. Ludvigson, and S. V. Nieuwerburgh (2017): "The Macroeconomic Effects of Housing Wealth, Housing Finance, and Limited Risk Sharing in General Equilibrium," *Journal of Political Economy*, 125, 140–223.

Favilukis, J., P. Mabille, and S. Van Nieuwerburgh (2022): "Affordable Housing and City Welfare," *The Review of Economic Studies*, 90, 293–330.

Ferrero, A., R. Harrison, and B. Nelson (2023): "House Price Dynamics, Optimal LTV Limits and the Liquidity Trap," *The Review of Economic Studies*, rdad040.

Garriga, C. and A. Hedlund (2020): "Mortgage Debt, Consumption, and Illiquid Housing Markets in the Great Recession," *American Economic Review*, 110, 1603–1634.

Greenwald, D. L. and A. Guren (2021): "Do Credit Conditions Move House Prices?" *NBER Working Paper Series*, w29391.

Iskhakov, F., T. H. Jørgensen, J. Rust, and B. Schjerning (2017): "The endogenous grid method for discrete-continuous dynamic choice models with (or without) taste shocks," *Quantitative Economics*, 8, 317–365.

Justiniano, A., G. E. Primiceri, and A. Tambalotti (2019): "Credit Supply and the Housing Boom," *Journal of Political Economy*, 127, 1317–1350, publisher: The University of Chicago Press.

Kaplan, G., K. Mitman, and G. L. Violante (2020): "The Housing Boom and Bust: Model Meets Evidence," *Journal of Political Economy*, 128, 3285–3345.

Kelly, R., F. McCann, and C. O'Toole (2018): "Credit conditions, macroprudential policy and house prices," *Journal of Housing Economics*, 41, 153–167.

Kiyotaki, N., A. Michaelides, and K. Nikolov (2011): "Winners and losers in housing markets," *Journal of Money, Credit and Banking*, 43, 255–296.

Lambertini, L., C. Mendicino, and M. T. Punzi (2013): "Leaning against boom–bust cycles in credit and housing prices," *Journal of Economic dynamics and Control*, 37, 1500–1522.

Lyons, R. C. (2018): "Credit conditions and the housing price ratio: Evidence from Ireland's boom and bust," *Journal of Housing Economics*, 42, 84–96.

Muñoz, M. A. and F. Smets (2022): "Macroprudential policy and the role of institutional investors in housing markets," .

Nakajima, M. and I. A. Telyukova (2020): "Home equity in retirement," *International Economic Review*, 61, 573–616.

Paz-Pardo, G. (2021): "Homeownership and Portfolio Choice Over the Generations," *SSRN Electronic Journal*.

Pew Research (2021): "Who rents and who owns in the U.S." Accessed online at https://www.pewresearch.org/fact-tank/2021/08/02/as-national-eviction-ban-expires-a-look-at-who-rents-and-who-owns-in-the-u-s/.

Piazzesi, M. and M. Schneider (2016): "Housing and macroeconomics," *Handbook of Macroeconomics*, 2, 1547–1640.

Queiró, L. and J. Oliveira (2022): "Mortgage Borrowing Caps: Leverage, Default, and Welfare," Working paper.

Sommer, K. and P. Sullivan (2018): "Implications of US tax policy for house prices, rents, and homeownership," *American Economic Review*, 108, 241–74.

Van Bekkum, S., R. M. Irani, M. Gabarro, and J. L. Peydró (2019): "Macroprudential policy and household leverage: Evidence from administrative household-level data," .

Appendix

This Appendix is organized as follows. Section A describes alternative models which differ in their assumptions about the housing market. These additional models are useful to validate the robustness of our results to changes in these key assumptions. Section B analyzes how sensitive our results are to various choices within the baseline model. Section C outlines the algorithms used for the computation of the equilibrium. Section D contains some additional plots, while section E presents additional data facts that support our modeling choices.

A. Alternative models

A.1. Model with a single quality

In this version of the model, we assume that there is only one house quality, which directly implies that the the aggregate amount of housing does not change. In fact, one can normalize it to unity, i.e. $H = 1$. In this case, it is possible to find the equilibrium house price analytically through the housing producer's first order condition:

$$p_h(1-\alpha)A_h\left(\frac{\bar{L}}{S}\right)^{\alpha} - 1 = 0 \;\Leftrightarrow\; p_h = \frac{1}{A_h(1-\alpha)}\left(\frac{S}{\bar{L}}\right)^{\alpha} \tag{19}$$

and the housing market clearing condition,

$$A_h\bar{L}^{\alpha}S^{1-\alpha} = \delta_h H \;\Leftrightarrow\; S = \left(\frac{1}{A_h}\delta_h\bar{L}^{-\alpha}\right)^{\frac{1}{1-\alpha}} \tag{20}$$

since after substituting equation (20) into (19), the per unit house price is simply a function of model parameters α, A_h and δ_h, as well as the amount of buildable land used in production, which is also fixed by assumption.

Importantly, note that in this case a change in mortgage credit conditions, i.e. changes in λ_{LTV} and λ_{LTI}, won't affect the equilibrium house price, p_h. In our model, however, this does not imply that the relative price of housing versus that of renting does not change as we move credit conditions. In fact, it does! As shown in Table 5 below, the rent to price ratio increases by 5.98%. Nevertheless, this change corresponds entirely to the increase in rental prices associated to the higher rental demand. In other words, the change in the ownership composition of the housing stock, which moves from homeowners to landlords, does not affect house prices because owner-occupied and buy to let properties are assumed to be identical, i.e. there is no segmentation.

This is still an important result because we are able to generate movements in the rent to house price ratio in response to a change in credit conditions even without the need to impose some degree of

Table 5: Effects of credit tightening under different assumptions on housing supply

	Single quality	Quality Ladder	Two Techn.
Δ Rent-to-Price	5.98 %	8.76 %	23.58 %
Δ Average house price to income	0.0 %	-0.65 %	-0.58 %
Δ Owner-occupied house price to income	0.0 %	-0.14 %	-2.14 %
Δ Buy-to-let house price to income	0.0 %	-0.14 %	8.11 %
Δ Rent to Income	5.98 %	8.06 %	22.86 %
Δ Homeownership rate	-2.34 p.p	-2.79 p.p	-2.88 p.p.
Δ Share of households with 3 properties	0.58 p.p.	0.68 p.p.	1.01 p.p.
Δ Share of houses in hands of 3-property landlords	0.81 p.p	1.24 p.p.	-4.42 p.p.

segmentation between the two markets. The key element of the model that generates this behavior is the presence of small heterogenous landlords that weight costs and benefits before buying an additional house and rent it out. In particular, this feature generates an upward sloping rental supply curve (see Figure 3) which is key in generating changes in both prices and quantities as credit tightens.

A.2. Model with a quality ladder: the baseline

In this version of the model we assume that houses differ in their quality or size. For simplicity, we further assume that there are only two types: owner-occupied (OO) and buy-to-let (BTL) with the former being larger or of better quality than the latter, i.e. $\tilde{h}_{oo} > \tilde{h}_{btl}$. We calibrate these two parameters to match the owner-occupied to rental house price ratio in the data.

An implication of this assumption is that the aggregate amount of housing is now endogenous and in fact given by: $H = \tilde{h}_{oo}H_{oo} + \tilde{h}_{btl}(1 - H_{oo})$ where H_{oo} also denotes the homeownership rate. As a result movements in credit conditions also generate changes in house prices through a composition effect. Since a credit tightening reduces the homeownership rate, the total aggregate amount of housing also drops because there are fewer high quality owner-occupied houses relative to lower quality rental houses. This drop in aggregate housing affects the per unit house price through equation (21) below

$$p_h = \frac{1}{1-\alpha}\left(\frac{1}{A_h}\right)^{\frac{1}{1-\alpha}}\left(\frac{\delta_h H}{\bar{L}}\right)^{\frac{\alpha}{1-\alpha}} \tag{21}$$

which we have obtained after substituting the housing investment function (4) into equilibrium condition that states that all depreciated houses must be refurbished: $Y_h = \delta_h H + (1/T)\tilde{h}_{oo}$. Equilibrium condition (4) is obtained by solving for S in the housing good producer first order condition and then plug it into the housing technology $Y_h = A_h \bar{L}^{\alpha} S^{1-\alpha}$.

In short, in a model with landlord heterogeneity, allowing for a quality ladder in the housing market is a way of generating changes in both house prices and rents in response to a credit shock, as shown in the second column of Table 5.

A.3. Model with two housing technologies

In the baseline model, the housing good producer uses the same production function to build both owner-occupied and buy to let properties. Alternatively, one could also assume that this is not the case since these are two completely different goods (full segmentation). Hence, in this version of the model, we assume that both owner-occupied and buy-to-let housing producers are independent of each other. Nevertheless, we further assume that they use the same Cobb-Douglas technology but with different degrees of efficiency. As a result, the equilibrium house prices for each housing sector $j = \{oo, btl\}$ are given by:

$$p_h^j = \frac{1}{(1-\alpha)} \left(\frac{1}{A_h^j}\right)^{\frac{1}{1-\alpha}} (\delta_h H_j)^{\frac{\alpha}{1-\alpha}} \tag{22}$$

where now the efficiency units in production is industry specific, A_h^j, and the total stock of housing is split between the two industries, i.e. $H = H_{oo} + H_{btl}$.

In this specification, the values of A_h^{oo} and A_h^{btl} are key to determine the price relationship between owner-occupied and buy-to-let properties. Therefore, we calibrate these two values such that the equilibrium house price ratio in the pre-reform economy equals to that in the data.[2] In fact, it is easy to show that:

$$\frac{A_h^{oo}}{A_h^{btl}} = \left(\frac{p_h^{btl}}{p_h^{oo}}\right)^{1-\alpha} \left(\frac{H_{oo}}{1-H_{oo}}\right)^{\alpha} \tag{23}$$

which allows us to pin down their ratio directly from the data, conditional on a specific value of the share of buildable land used in production.[3] In any case, it is important to note that house prices in this economy differ not only because the efficiency parameters do not coincide, but also because the share of each housing good in the total housing stock differs.

Bearing that in mind, we now turn to the results of our main experiment. Under these assumptions, a tightening of the credit conditions also generates an increase in the rent to house price ratio (23.58%) as well as a reduction in the homeownership rate (-2.88 p.p.). Unlike in the baseline, owner-occupied and buy-to-let property prices move in opposite directions since the drop in homeownership rate leads to a reduction in owner-occupied house prices, while the increase in the rental share induces an increase in buy-to-let property prices. As a result, the effect on average house prices will depend upon the strength of each of these two channels and the equilibrium quantities of the two housing goods. For the model parametrization considered, and explain below, owner-occupied house prices drop by -2.14%, while buy-to-let house prices rise by 8.11%. However, since there are a lot more owner-occupied houses in the economy (approx. 80%), the average house price to income ratio slightly decrease (-0.58%) in response to a credit tightening.

[2] On average owner-occupied housing is 1.21 times more expensive than buy-to-let properties.

[3] Note that H_{oo} corresponds to the home-ownership rate since every household must buy an owner-occupied property before purchasing additional buy-to-let properties. Thus, the two ratios in the RHS of equation (23) are observable.

A.4. Parameterizing the alternative models

We calibrate the alternative models to have a fairer comparison of the effects of the reform among the different versions of our model economy. In particular, we choose the same three parameters: discount factor, β, utility premium from homeownership, $s(h)$, and the scaling factor in housing production, A_h, to jointly match the wealth to income ratio, the homeownership rate, the house price to income ratio and the house price to rent ratio.[4]

Table 6 reports the model fit under each of the three assumptions. The three versions of the model under-predict the wealth to income ratio. Nonetheless, the model with one single quality is closer to this target in part due to higher levels of housing wealth. In this set up both the homeownership rate and the average house price to income ratio are slightly above target, while they are below target for the models with two different house prices. The house price to income ratio is matched reasonably well in all three versions of the model.

In terms of untargeted moments, the share of households with three properties is under-predicted in all three versions since we cap the maximum number of houses owned by a household to three. In any case, the model with a single quality predicts a larger share of households with three properties, which may be explained by the higher equilibrium rental price in this economy as it incentives saving in housing instead of in bonds. Nevertheless, the rent to income ratio is also accurately predicted by the other two versions of the model. Finally, all the three versions of the model are also able to replicate average housing status along the age distribution despite not reported here.

Table 6: Targeted and untargeted moments under different model assumptions

Moment	Data	Single quality	Quality ladder	Two Tech.
Targeted:				
Wealth to income ratio	6.78	6.00	5.32	4.96
Homeownership rate	80.00%	80.14%	79.13%	79.41%
Average house price to income ratio	5.00	5.23	4.90	4.78
House price to rents ratio	22.58	22.56	23.00	22.54
Untargeted:				
Rents to income ratio	0.2216	0.2319	0.2132	0.2124
Share of households with 3+ properties	5.11%	4.18%	3.97%	3.52%

[4]Note that in the model with two housing technologies we pick A_h^{oo} and let A_h^{btl} be determined by the relationship in (23).

B. Robustness

B.1. Aggregate housing supply elasticity

We have discussed in Appendix A.2 that in the baseline model house prices change in response to a credit tightening because the amount of aggregate housing varies endogenously with the equilibrium homeownership rate. Moreover, we also know from the market clearing conditions that the amount of aggregate housing is proportional to the housing production. Thus, in this section we study how the elasticity of aggregate housing supply affects our results.

This elasticity is governed by the share of land used in production α. To see this, compute the elasticity of housing production with respect to the per unit house price:

$$\begin{aligned} \frac{\partial Y_h}{\partial p_h}\frac{p_h}{Y_h} &= A_h^{1/\alpha}\bar{L}\frac{1-\alpha}{\alpha}\left((1-\alpha)\,p_h\right)^{\frac{1-\alpha}{\alpha}-1}\frac{p_h}{Y_h} \\ &= \frac{1-\alpha}{\alpha}\left((1-\alpha)\,p_h\right)^{-1}p_h \\ &= \frac{1}{\alpha} \end{aligned} \tag{24}$$

Therefore, we modify the value of α and recompute the effect of the reform across steady states. In doing so, we make sure that the house price in the pre-reform economy does not change and modify the efficiency of housing production A_h accordingly. Figure 8 shows the changes in the average house price, the homeownership rate and the rental rate across steady states for various elasticity levels.

Figure 8: Effects of the reform as a function of the supply elasticity

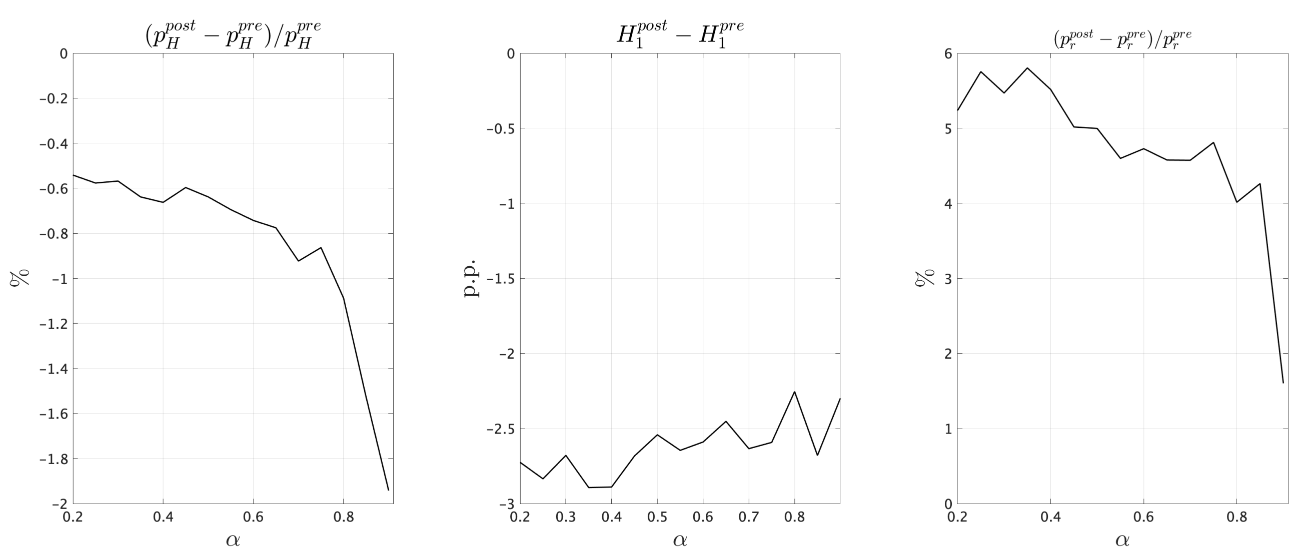

(a) Average house price (b) Homeownership (c) Rental prices

Recall that a credit tightening forces some households to postpone or cancel their buying decisions which in turn leads to a higher share of renters in the economy. Remember as well that in the model with a quality ladder, movements in the share of renters are the only driver of movements in average house prices. It is purely a composition effect: less high quality owner occupied houses relative to low quality buy-to-let properties implies a fall in average house prices. Therefore, for large values of α, i.e. for economies in which housing investment is less sensitive to prices changes, one would need a bigger drop in average house prices to accommodate the same change in the housing stock. In other words, if the housing producer does not react as much to changes in credit conditions then house prices have to drop more to clear the markets – see panel (a). In panel (b), we show that the change in homeownership rate is higher the higher the elasticity of housing supply, even though the slope is pretty flat. Finally, panel (c) shows that rental prices rise more in economies with high aggregate housing supply elasticity. This is partly explained by the lower fall in house prices which doesn't allow many new landlords to step in and provide additional rental houses. As a result, rental prices need to increase more to clear the rental market. Moreover, the larger drop in the homeownership rate in the economies with high aggregate housing supply elasticity also puts more pressure into the rental market: there is a larger share of households competing for a rental home since they cannot afford to buy their own at the new limits, which combined with the small changes in rental supply puts upward pressure on rental prices.

Finally, we also compute the consumption equivalent variation (CEV) across steady states for a new born with zero assets and zero housing wealth. The solid black line in Figure 9 depicts this average CEV for economies with different housing supply elasticity. As expected, the welfare loss is more severe in economies with high housing supply elasticity (low α) because the rise in rental prices is larger, the fall in house prices is tiny and there is a large drop in homeownership.

Figure 9: CEV – the role of housing supply elasticity

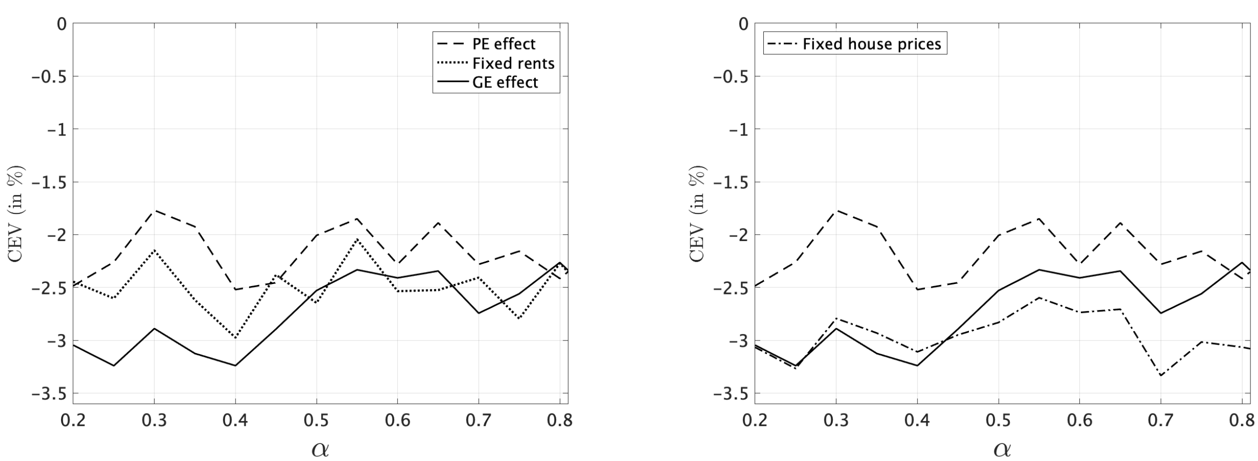

(a) House price decomposition

(b) Rental price decomposition

We also compute these welfare numbers in an economy with fixed rental and housing prices (dashed line) and show that, as one would have expected, the CEV is constant across different values of the housing supply elasticity. An obvious implication of this result is that the response of real estate prices to a credit tightening, which in itself depends on the housing supply elasticity, is essential in assessing who are the ultimate winners and losers from the reform.

Finally, since we are interested on how much of these differences in welfare effects can be attributed to changes in rents and how much corresponds to house prices, we also compute the CEV under fixed rental prices (dotted line) and fixed house prices (dash dotted line). Starting from panel (a) and comparing the solid line and the dotted line we see that the gap between these two lines is wider for low values of α. In this region, the rise in rental prices is the largest making welfare losses more severe. Moving to panel (b) and comparing the dashed dotted line withe solid line, we can observe that the gap between these two lines is wide for high values of α. In this case, however, the solid line is on top since the larger fall in house prices when the housing supply is very elastic makes the effects of the reform less harmful.

C. Solution method

The solution of the model consists of two main loops: an inner loop that solves the household problem given structural parameters and prices, and an outer loop that recovers the equilibrium distribution and prices. In what follows, we describe each of them in detail.

C.1. General equilibrium

Algorithm 1 explains how to recover the equilibrium house price p_h^* and rental rates p_r^*, as well as the equilibrium distribution of households in the model economy, $\mu^*(\cdot)$. It works as follows. First, we start by guessing a value for the rental price and the homeownership rate. Note that the latter allows us to compute the per unit price of housing analytically through the market clearing conditions. Given house and rental

---

Algorithm 1: General Equiliibrium

Input:

Vector of parameters: Θ

Rental price guess: p_r^{guess}

Homeownership rate guess: HO^{guess}

Output:

Equilibrium rental rate: p_r

Equilibrium avg. house price: p_h^{avg}

Stationary distribution: $\mu(\cdot)$

$p_r^1 \leftarrow p_r^{guess}, HO^1 \leftarrow HO^{guess}$;

$p_h^1 \leftarrow HO^1$;

for i = *1 : max iter* **do**

  $\phi^a(\cdot), \phi^h(\cdot) \leftarrow$ Algorithm 2: Household problem;

  $\mu(\cdot) \leftarrow \phi^a(\cdot), \phi^h(\cdot), \mu_0$;

  $R^d, R^s \leftarrow \mu(\cdot)$;

  if $\max\left\{abs\left(R^d - R^s\right), abs\left(HO^i - (1 - R^d)\right)\right\} < tol$ **then**

    return $p_r^i, p_h^i, \mu(\cdot)$

  else

    if $R^d > R^s$ **then**

      $p_r^{i+1} > p_r^i$

    else

      $p_r^{i+1} < p_r^i$

    end

    $HO^{i+1} \leftarrow \xi HO^i + (1 - \xi)(1 - R^d)$;

  end

end

$p_h^{avg} \leftarrow (1 - R^d)p^h(\tilde{h}^1) + R^d p^h(\tilde{h}^2)$;

---

prices, it is possible to solve the household problem for which we use the algorithm described below. As a result, we obtain the asset and housing tenure policies, which are probabilistic distributions since we use extreme value I shocks. These policies are then used to back up the stationary distribution over asset, income, housing tenure and age. In the next step, we use the stationary distribution to compute the share of renters, homeowners, and landlords in our model economy. Note that the share of renters corresponds to the demand of rental housing, while the supply can be computed using the shared of landlords adjusted by the number of properties that each of them rents out. Given these two quantities, we check if their difference is below a predetermined tolerance level. If that is the case, we exit the algorithm since we have found an approximate equilibrium. Otherwise, we update our guesses for the rental price and the homeownership and start over again from the household problem.

D. Additional model-generated figures

Figure 10: Transition Dynamics – some extra variables

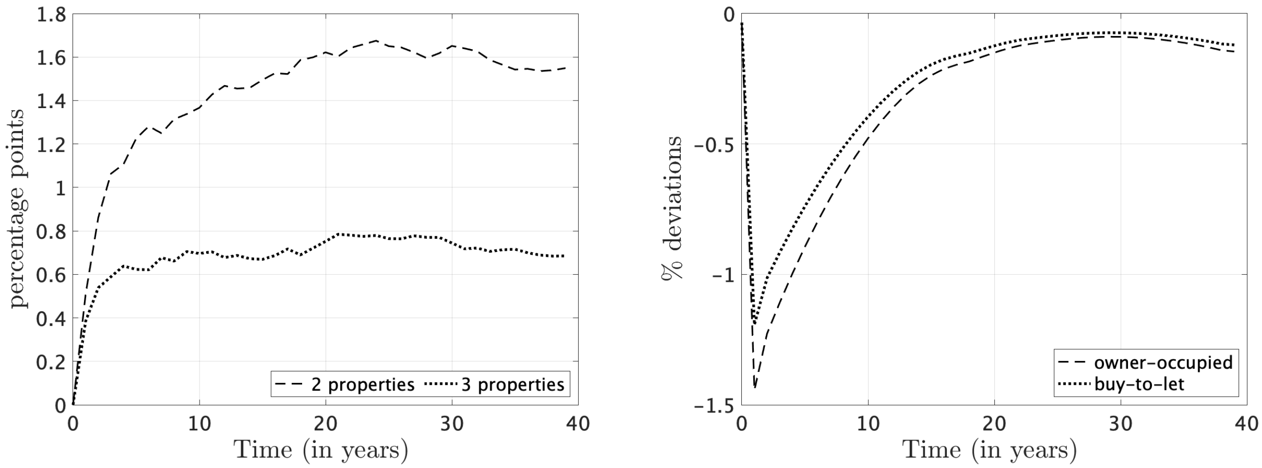

(a) Share of landlords

(b) House Prices

Figure 11: An alternative decomposition of the CEV – the role of house prices

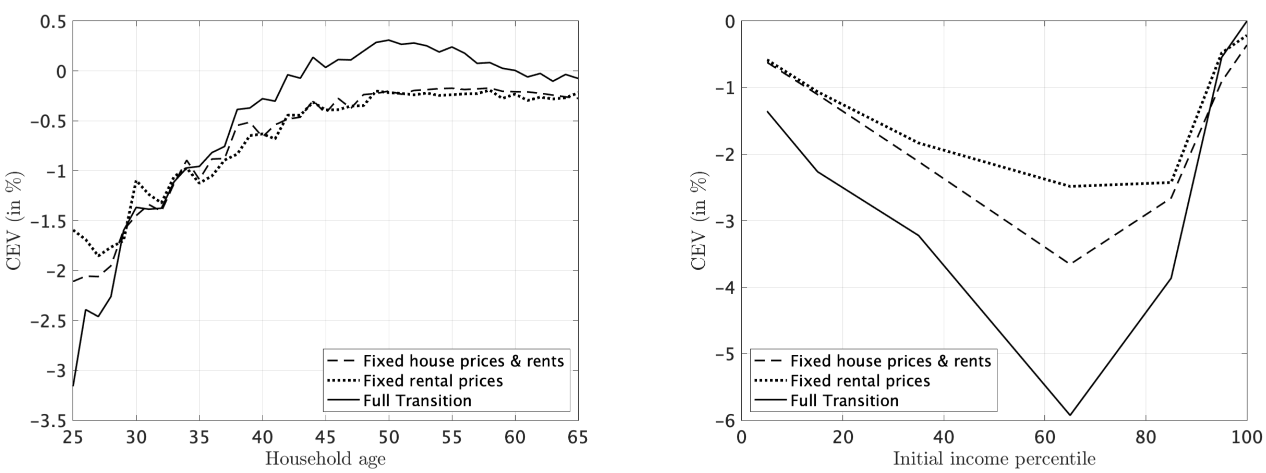

(a) Age

(b) Income

E. Additional data facts

E.1. Structure of the Irish rental sector

In the Irish rental sector, the vast majority of houses to rent are owned by relatively small landlords. Figure 12, right panel, shows that around 37% of all rental properties are owned by people that only own that property, and the next 15% is owned by people that only own two. In contrast, big investors with more than 50 properties represent less than 10 percent of all rentals.

Figure 12: Irish rental sector structure

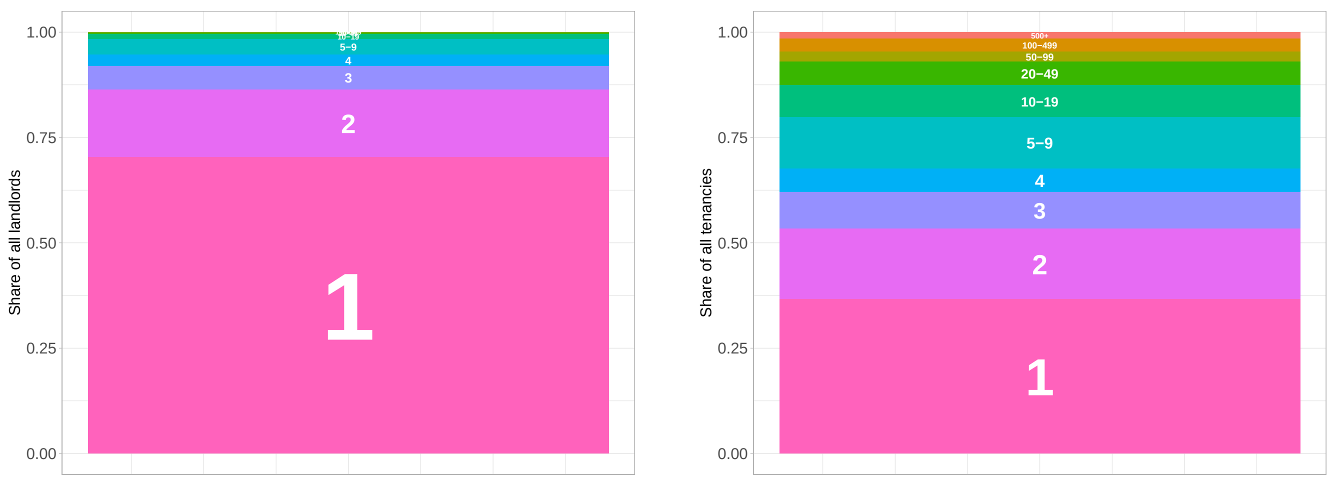

(a) Share of landlords by number of properties owned

(b) Share of all tenancies owned by landlords, depending on the number of properties they own

Source: Department of Finance (2019)

Despite their small aggregate size, non-household and other rich investors have grown over time, and concentrated in particular housing markets, such as newer developments around Dublin and other cities. In Figure 13 we show that, although they are still a small share of all transactions (left hand side) they do buy a significant amount of non-owner-occupied housing (right hand side). However, around the time of the reform around 60% of transacted non-owner-occupied housing was bought by household buyers.

Figure 13: Share of property transactions, by type of buyer and year.

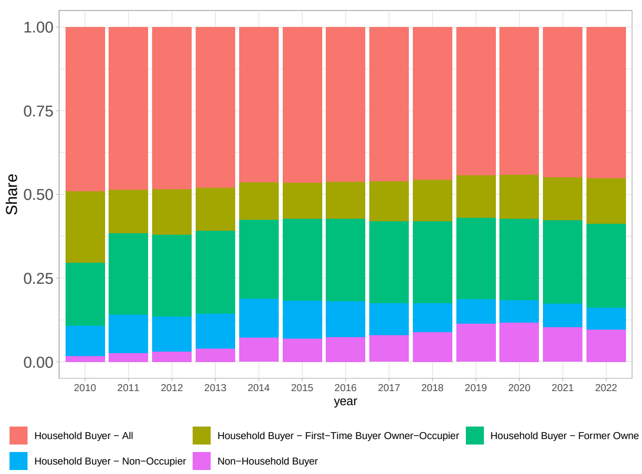

(a) All properties

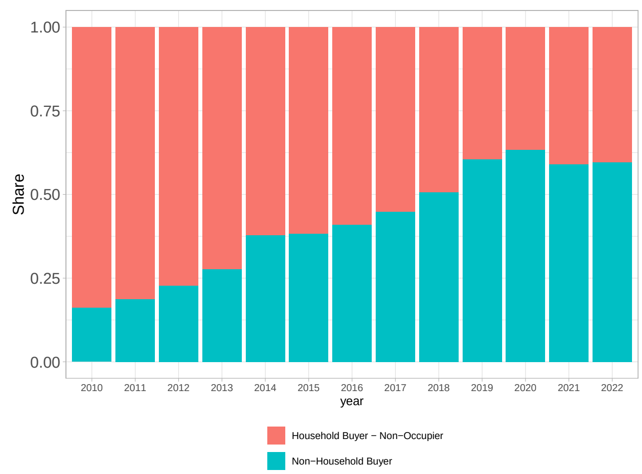

(b) All properties bought by non-owner-occupiers

Source: Irish Central Statistics Office.