Online Appendix

Appendix A. Additional empirical evidence

A.1. Irish rental sector

In our model economy we assume that the rental sector is populated by households that own one or two rental properties. Although, this assumption may seem restrictive, it is consistent with the Irish private rental sector.

In Figure A1, we use data from the Central Statistical Office (CSO) on residential property transactions to show that the vast majority of non-occupier property purchases correspond to household buyers. In fact, in 2015, the year when the macroprudential reform was introduced, around 70% of those transactions correspond to household buyers. Nonetheless, these data also confirms that the role of non-household buyers such as pension funds, private rental firms and Real Estate Investment Trusts (REITs) has increased over the last decade.

In Figure A2, we dig deeper into the ownership structure in the rental sector and use data from the Residential Tenancies Board (RTB) using RTB registrations as a proxy for ownership. Panel A shows the share of landlords by number of tenancies. Note that a tenancy is not fully analogous to a property as there may be some instances where

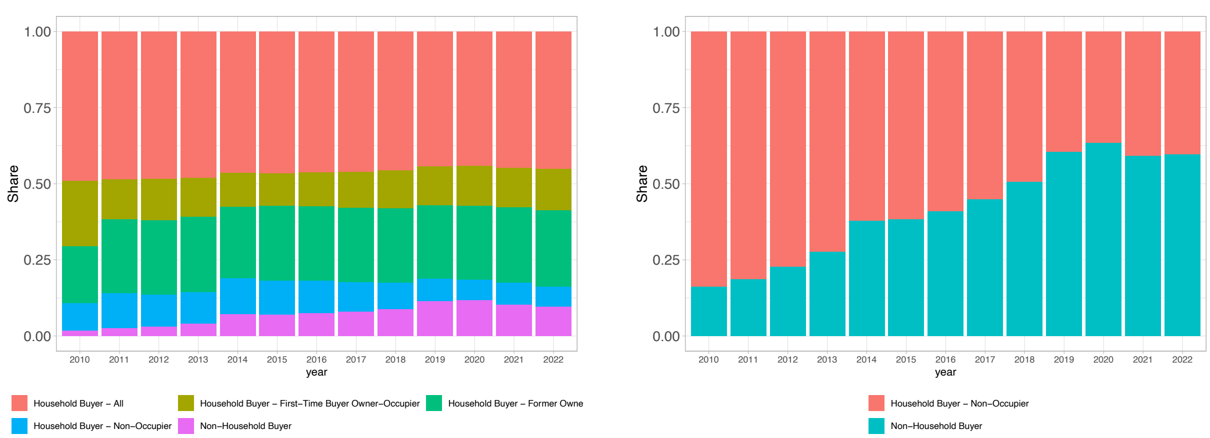

A. All properties B. All properties bought by non-owner-occupiers

FIGURE A1. Share of property transactions, by type of buyer and year

NOTE: This figure shows the share of all house sales by type of buyer and year in panel A. Panel B focus the attention in non-occupier buyers which are split into two categories: household buyers and non-household buyers. Data is available at the CSO.

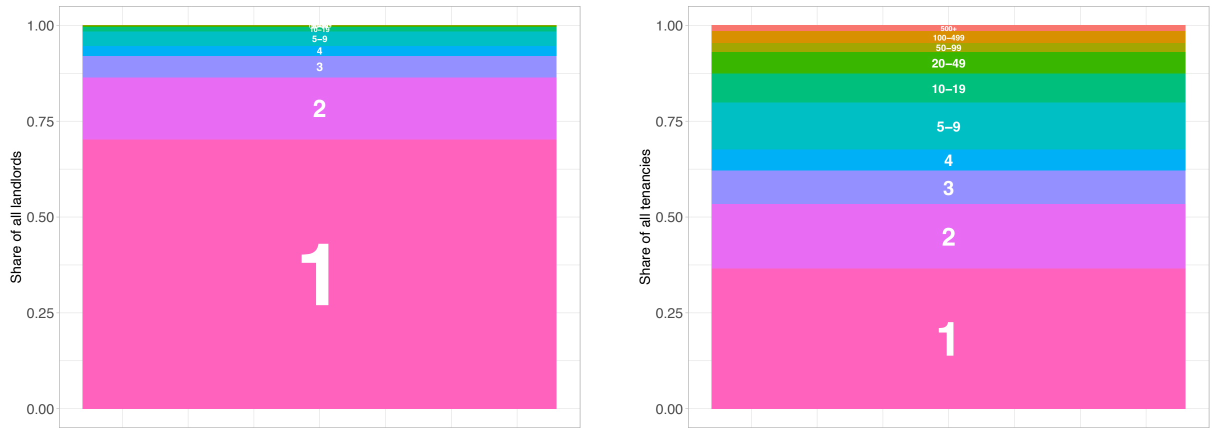

A. Share of landlords by number of tenancies

B. Share of all tenancies owned by landlords

FIGURE A2. Irish rental sector structure

NOTE: This figure shows the share of landlords by number of registered tenancies (panel A) as well as the share of tenancies owned by landlords with different number of registered tenancies (panel B). Data is from the RTB.

there are multiple tenancies in one property (e.g. a flat with multiple rented rooms). Nonetheless, the vast majority of tenancies are individual properties. With that in mind, the evidence on the RTB data points to a lesser role of large scale professional landlords as only 4.6% of tenancies are held by landlords with more than 100 units. On the other hand, the vast majority of landlords register a single rental property (70%) or at most two (86%). One gets a similar picture, if looks at the share of tenancies by landlords – panel B. In fact, landlords with one or two properties registered more than 50% of all tenancies.

Figure A1 and A2 are consistent with each other as the rise of institutional investors in Ireland is mostly concentrated in newly constructed, high quality and well located units, but not so relevant at the aggregate level (Ireland's Department of Finance 2019).

A.2. Macroprudential limits, house & rental prices

In this section, we describe the data used in our regression analysis, provide additional non-parametric evidence on the opposite response of house and rental prices to the introduction of macroprudential limits, and run some robustness test that verify such relationships.

A.2.1. Data sources

The core of our final data set is the result of combining the distance measure with county-level selling and rental house prices. In our main specifications, we borrow the distance measure from Acharya et al. (2022). They construct this measure using loan-level information on residential mortgages. In particular, they "calculate what would have been the distance from the limits for each borrower in the year before the policy, assuming that the limits were in place during that period" (p. 12, Acharya et al., 2022). For confidential reasons, we got this information aggregated to the county level.

Data on house and rental prices comes from Daft.ie. We borrow these data from Lyons (2018) since in his website he has the aggregated time series for each Irish county of both selling and rental prices.

A.2.2. Non-parametric evidence

Figure A3 shows the variation in house price growth (panel A), the distance measure (panel B) and rental price growth (panel C) across all Irish counties. In low-distance counties, such as those areas around Dublin, house price growth was slower and close to zero, while rental prices were growing faster at a pace around 30-35%. This observation suggests that the distance measure is positively correlated with house price growth while it is negatively correlated with rental price growth. This statement was formally verified in our regression analysis in Section 4.2.

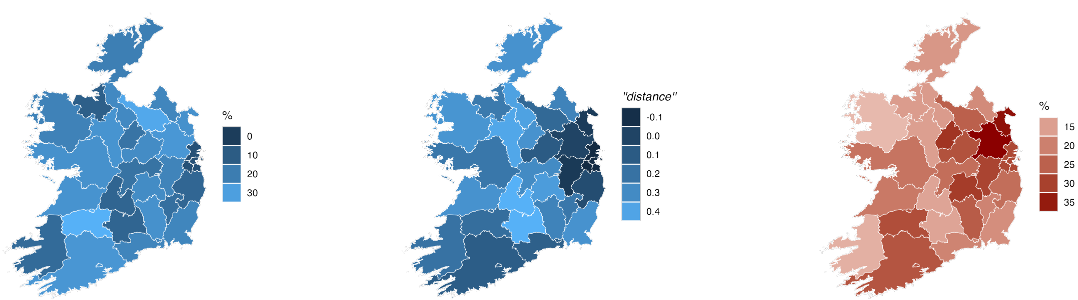

A. House price growth B. Distance measure C. Rental price growth

FIGURE A3. Counties, lending limits, house & rental price growth

NOTE. This figure shows the county-level distance from the limits (panel A), house price (panel B) and rental price (panel C) growth between the third quarter of 2014 and the fourth quarter of 2016. Data on prices comes from Daft.ie while the distance measure was provided by Mateo Crosignani and corresponds to the one in their paper: Acharya et al. (2022). Darker colors indicate less distant counties, lower house price growth and higher rental price growth.

Appendix B. Further model details

B.1. Solution method

The steady state solutions of the model consist 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. A description of the algorithms used for the approximation of the steady state equilibria can be found in Appendix B.1.1 and B.1.2. In addition to steady state equilibria, welfare comparisons also require to solve the transition from one steady state to another. The computational approach used to solve for such transition is described in Appendix B.1.3.

B.1.1. Household problem

As shown in Section 2.1, the household state variables are age, j, income, y, the housing state, $s = (h, \tilde{h})$, and net financial wealth, a. Consequently, the first step is to discretize the continuous state variables. Financial wealth lie on a non-linearly spaced grid with 150 points that includes 50 negative values and 100 positive ones, while the stochastic component of income is discretized using the approach in De Nardi, Fella, and Paz-Pardo (2020) that accounts for non-linearities and age-dependence. In particular, we allow for 7 points for the stochastic component of income whose values vary with the working age of the household.[20] The remaining state variables are already discrete. The model period is one year, households live up to 71 years, and the housing state can take 7 different values: *(i)* $s = (0, \tilde{h}_1)$ if renter, *(ii)* $s = (1, \tilde{h}_2)$ if small owner, *(iii)* $s = (1, \tilde{h}_3)$ if mid owner, *(iv)* $s = (1, \tilde{h}_4)$ if big owner, *(v)* $s = (2, \{\tilde{h}_1, \tilde{h}_3\})$ if mid landlord with one rented house, *(vi)* $s = (2, \{\tilde{h}_1, \tilde{h}_4\})$ if big landlord with one rented house or *(vi)* $s = (3, \{2 \times \tilde{h}_1, \tilde{h}_4\})$ if big landlord with two rented houses.

Since households die with certainty at age J, we know their optimal policy in their terminal period, so we can proceed by backward induction and compute the remaining age-dependent policy functions. Note that households make the standard consumption-savings choice, a', as well as decide on the next period housing tenure, s', at each age. Given that the housing choice is discrete, the solution of the household problem requires using computational techniques employed to solve discrete-continuous dynamic choice models. We follow closely the recipe from Fella (2014) and Iskhakov et al. (2017) to use

[20] The transition matrix that controls the evolution of household's income over time is also age-dependent and hence it is of dimension $7 \times 7 \times 41$ where 41 is the retirement age J^{ret}.

the endogenous grid method (EGM) together with taste shocks to solve for these discrete-choice specific policy and value functions. In a nutshell, for each $j < J$ we first compute the expected marginal utility to then invert the Euler equation and get the endogenous consumption-asset policy in a normal EGM step. After that, we apply the general EGM procedure to verify the global optimality of these choices in the non-concave region and discard those that are not. Finally, we use the obtained s'-dependent value and policy functions to compute the probability of the discrete choice using the Logit probability formula and the expected value function using the log-sum formula. These are stored and used in the next step of the backward induction. Once the backward induction is finished, the final outcomes of the algorithm are s'-dependent consumption-savings policy functions, a discrete choice probability and a value function.

B.1.2. General equilibrium

To compute the equilibrium in the housing and rental markets we proceed as follows:

1. Make a guess for the rental price, p_r^g, and the per unit house price p_h^g.
2. Use the per unit house price guess to analytically recover the quantity of housing produced in the construction sector using equilibrium relationship (12). Then, use the steady state version of the housing law of motion, equation (15), to obtain the level of housing stock implied by the construction firm's optimal choices given the per unit house price guess: H_s^g.
3. The implied price of land p_L can at this point also be recovered as it is an exponential function of the per-unit housing price p_h. Note that after substituting the housing investment function (12) into the first order condition with respect to land, one would get:

$$p_{L,t} = \alpha_L (1 - \alpha_L)^{\frac{1-\alpha_L}{\alpha_L}} \left(p_{h,t} A_h \right)^{\frac{1}{\alpha_L}}$$

4. Given rental and per-unit house price guesses $\left\{ p_r^g, p_h^g \right\}$, compute the transacted house prices for each quality, $p^g(\tilde{h}_n) = h_n p_h$. Then, use the algorithm described in Appendix B.1.1 to get the value and policy functions that solve the household problem.

5. Using the household's consumption-saving policy and the discrete choice probability, recover the stationary distribution of households $\mathcal{D}(a, s, y, j)$ as it contains all the information needed for evaluating if the rental and housing market clear.

   a. Rental demand equals the share of households that choose to be renters:

$$R^d = \sum_{i_a=1}^{n_a} \sum_{i_y=1}^{n_y} \sum_{j=1}^{J} \mathcal{D}(a_{i_a}, s_1, y_{i_y}, j)$$

   b. Rental supply is given by the sum of landlords with one rented out property plus two times the share of landlords with two rented out properties:

$$R^s = \sum_{i_a=1}^{n_a} \sum_{i_s=5}^{6} \sum_{i_y=1}^{n_y} \sum_{j=1}^{J} \Big(\mathcal{D}(a_{i_a}, s_{i_s}, y_{i_y}, j)\Big) + 2 \times \sum_{i_a=1}^{n_a} \sum_{i_y=1}^{n_y} \sum_{j=1}^{J} \mathcal{D}(a_{i_a}, s_7, y_{i_y}, j)$$

   c. The aggregate housing stock from the demand side is a quality weighted (row vector) sum of the share of households who live in each type of home (column vector):

$$H_d^g = \begin{bmatrix} \tilde{h}_1 & \tilde{h}_2 & \tilde{h}_3 & \tilde{h}_4 & \tilde{h}_3 & \tilde{h}_4 & \tilde{h}_4 \end{bmatrix}_{1\times 7} \times \left[\sum_{i_a=1}^{n_a} \sum_{i_y=1}^{n_y} \sum_{j=1}^{J} \mathcal{D}(a_{i_a}, :, y_{i_y}, j) \right]_{7\times 1}$$

6. If $|R^d - R^s| < \varepsilon_r$ and $|H_s^g - H_d^s| < \varepsilon_h$, then we are done. Otherwise, we need to update the guesses and go back to step 2.

The final outcomes of this algorithm are: an equilibrium rental price, an equilibrium per-unit and average house price, an implied land price, the stationary distribution of households over their state space and optimal policy and value functions.

B.1.3. Transition dynamics

To compute the transition paths shown in Figures 6 and 8, we resort to the traditional approach that assumes that at time $t = 0$ the economy is initially in a steady state. Then, at $t = 1$ the policy reform is introduced as a surprise for households and maintained forever. Recall that in the macroprudential experiment the policy reform consists in introducing tighter LTV and LTI limits, so that λ_{LTV} and λ_{LTI} change while everything else remains untouched; while for the interest rate experiment, it is only the return on financial assets r_s and the mortgage rate r_b that increase in the new steady state. Note that for the transitory shock this new steady state is identical to the initial steady state, while the exogenous path of interest rates changes. In either case, the key idea is to assume that after T periods the transition from the old to the new steady state is completed. As a result, one can safely assume that policy and value functions at time $t = T$ are those from the new steady state. So that $c_T = c_{ss}^{new}$, $a_T = a_{ss}^{new}$ and $\mathbb{P}_T(s) = \mathbb{P}_{ss}^{new}(s)$.

For a given sequence of prices $\left\{ p_t^r,\ p_t(\tilde{h}_n) \right\}_{t=1}^{T}$, the previous insight allow us to solve the household problem backwards and obtain their policy functions at each point in time $\left\{c_t, a_t, \mathbb{P}_t(s)\right\}_{t=1}^{T}$. Knowing that $\mathcal{D}_0 = \mathcal{D}_1$, these are useful to iterate the distribution forward: $\mathcal{D}_{t+1} = \Gamma_t\left(\mathcal{D}_t\right)$ where Γ_t is the mapping obtained from the policy functions. Finally, using the sequence of household distributions over their state space $\{\mathcal{D}_t\}_{t=0}^{T}$ one can check if rental and housing markets clear at each point in time. If they do not, then the given sequence of prices needs to be updated until they do.

Thus, the most difficult aspect of the transition is to find suitable paths for rental and house prices. We approach this problem by first guessing different rental and housing price paths and evaluating ex-post which ones are closer to form an equilibrium sequence in housing and rental markets. These guesses are constructed parametrically by imposing an initial jump and a degree of curvature in its reversal to the new steady state level. Once we have a sense on how these equilibrium paths should look like, we follow a similar approach to that described in point 6 of the general equilibrium algorithm with the caveat that we now update the guesses based on the gaps between supply and demand along the entire path and not just based on one point in time.

B.2. LTI and LTV implementation in Ireland

As stated in Section 2.1, the borrower must satisfy two constraints. First, a loan-to-income (LTI) requirement that limits household's borrowing to a multiple, λ_{LTI}, of its current (annual) income. And second, 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.

When Central Banks establish these limits, they often include some exemptions based on the type of borrower or the type of property households purchase. For example, the Irish reform of 2015 imposed a LTI limit of 3.5 that only applied to First Time Buyers (FTBs). In the model, we identify FTBs with households that transition from renting into owning as there are very few (or even zero) households that after selling their primary residence become homeowners for a second time during their life-cycle. For all other borrowers, we let the pre-reform limit to apply as this was the LTI implicitly imposed by banks in absence of the Bank of Ireland macroprudential framework. Hence, formally, the LTI in the *post-reform* economy is

$$a' \geq -\lambda_{LTI}^{post}\ y \qquad \text{if} \quad h' = 1 > h \tag{A1}$$

$$a' \geq -\lambda_{LTI}^{pre}\ y \qquad \text{otherwise} \tag{A2}$$

Moreover, the reform also included some exceptions for the LTV limit based on the type of purchase. For example, buy-to-let buyers faced a more stringent 70% loan-to-value limit. We include this feature in the model by distinguishing between owner-occupied and buy to let purchases for which we let $\lambda_{LTV}^{oo} = 0.8$ and $\lambda_{LTV}^{btl} = 0.7$ to apply. In the model, it is easy to identify this purchases as households that own more than one property always lease it out. Hence, in the *post-reform* economy the LTV limit is given by

$$a' \geq -\lambda_{LTV}^{oo}\ p(\tilde{h}_n') \qquad \text{if} \quad h' = 1, h = 0 \tag{A3}$$

$$a' \geq -\left(\lambda_{LTV}^{oo}\ p(\tilde{h}_n') + \lambda_{LTV}^{btl}\ p(\tilde{h}_1')\left(h' - 1\right)\right) \quad \text{if} \quad h' > 1 \geq h. \tag{A4}$$

B.3. Computing lifetime consumption equivalent variations

We evaluate the distributional effects of the 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 $x = (a, y, h, j)$, we compute the consumption equivalent variation $g(x)$ as

$$V_0(x; g) \equiv (1+g)^{1-\gamma} V_0(x) = V_1(x) \qquad \Rightarrow \qquad g(x) = \left[\frac{V_1(x)}{V_0(x)}\right]^{\frac{1}{1-\gamma}} - 1 \tag{A5}$$

where we are using the fact that the utility function is CRRA. From (A5) 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.

B.4. Estimation of the earnings process

Our earnings process is based on De Nardi, Fella, and Paz-Pardo (2020). Namely, we extract out the persistent and the transitory component of earnings using the procedure described in Arellano, Blundell, and Bonhomme (2017), and then incorporate the dynamics for the persistent component in a nonparametric way. 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, Fella, and Paz-Pardo (2020).

We use data from the Household Finance and Consumption Survey (HFCS) to extract the average age-earnings profile in the Irish economy after taking into account year effects. However, the triennial nature of the HFCS does not allow us to estimate an annual earnings process. Hence, for the stochastic component of the earnings process we use household earnings data for the United Kingdom from De Nardi, Fella, and Paz-Pardo (2024), who extract them from the BHPS/Understanding Society survey, and assume that the stochastic properties of household earnings are similar across both countries.

We have also estimated an annual earnings process for Ireland based on EU-SILC data (European Union Statistics on Income and Living Conditions), which, despite being nationally representative, is targeted to produce statistics on poverty and living conditions and hence might capture the earnings dynamics of the upper part of the income distribution in a more limited way. Our main results with this alternative earnings process are very similar.

Appendix C. Further model results

C.1. Supply-demand illustration over the transition path

The rental and house price responses to a credit tightening differs in the long and the short run as we describe in Sections 4.3 and 5. These differences are explained by a combination of factors that are embedded in the elasticities of rental and housing supply and demand.

Figure A4, as the transition counterpart of Figure 3, plots the supply and demand shifts in the period immediately after a temporary interest rate shock hits the economy. The key intuition from the supply and demand illustration when comparing steady states still holds and we see that as a result of the credit tightening there is an equilibrium in $t = 1$ that features higher rental prices (panel A), lower house prices (panel B) and higher rent-to-price ratio (panel C). However, the slope of the supply and demand curves both for rental and owner-occupied housing differs substantially. In particular, supply curves become much more inelastic in the short run. In the housing market, this is explained by the slow adjustment of the aggregate housing stock, which requires large short-run house price movements to induce the necessary adjustments in housing construction to clear this market (middle panel). In the rental market, it is the presence of fixed costs to become a landlord – arising from housing transaction costs and indivisibilities – which limits the responsiveness of rental supply in the very short run (left panel),

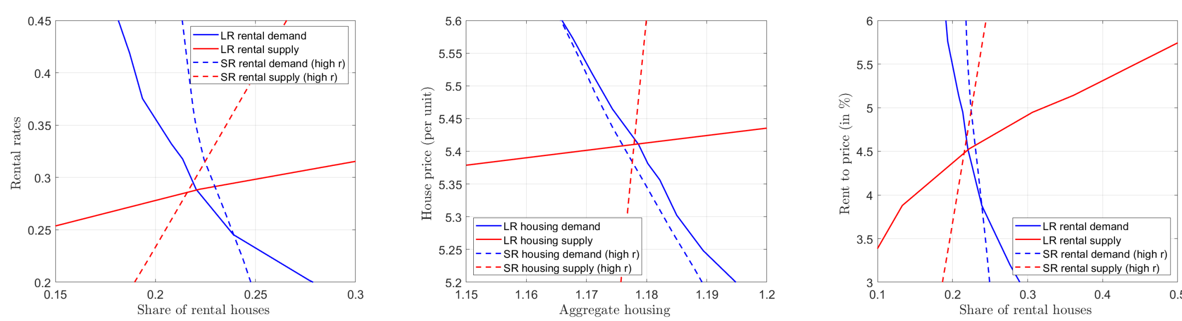

A. Rents B. House prices C. Rent-to-price ratio

FIGURE A4. Supply and demand in the rental and housing markets, baseline model

NOTE. This figures show the main mechanisms of the model through a supply and demand illustration. The demand and supply curves are computed numerically using a suitable parameterization of the model economy, and varying rental prices in the period after the shock hits while keeping house prices at their optimal transition path (panels A and C) and varying house prices after the shock hits (i.e. in $t = 1$) while rents stay at their optimal path (panel B).

relative to the medium run since households need time to adjust their saving decisions to overcome these barriers. As a result, both house prices and rents respond sharply to the credit tightening in the short run.

C.2. Interaction between mortgage credit limits

Loan to Value (LTV) and Loan to Income (LTI) limits are often introduced jointly, as it was the case for Ireland. However, there are some other countries, such as the United Kingdom that have only an LTI limit in place. To understand the contribution of each of them to the overall quantity and price effects of the reform, we compute two counterfactual economies: *(i)* the *Only LTI* economy which imposes the institutional 3.5 LTI limit but leaves the LTV unchanged, and *(ii)* the *Only LTV* economy which imposes the institutional 80% LTV limit but leaves the LTI unaltered.

TABLE A1. Non-linear interactions between credit limits

	Full-Reform	**Only LTI**	**Only LTV**
Δ% Rent-to-Price	+2.45 %	+2.46 %	+0.3 %
Δ% Rent-to-Income	+2.58 %	+2.48 %	+0.56 %
Δ% Price-to-Income	+0.13 %	+0.02 %	+0.23 %
Δ Homeownership rate	-3.18 p.p	-2.83 p.p.	-1.45 p.p.

NOTE. This table shows the effects of the reform on the rent to house price rate and home-ownership rate (first column) and decomposes the role of each limit by imposing one at a time. A tighter LTI (second column) has a larger effect than the tighter LTV (third column) if they are introduced on their own.

Table A1 shows the change relative to the *pre-reform* economy in the rent-to-price ratio as well as the homeownership rate after imposing both (full-reform) or one of these two limits. Results show that the LTI alone had a larger effect than the LTV, with a rise in the rent-to-price ratio similar to the full reform. The fall in the homeownership rate was also more pronounced in the counterfactual economy with only tighter LTI limits. Nonetheless, there are some interactions between LTV and LTI limits which shows here in the form of a larger drop in the homeownership rate for the full reform.[21] It is important to note, however, that these results are specific to our case study. In fact, the LTI introduced in Ireland was quite stringent relative to those introduced in other

[21] Other research has shown how these interactions can also affect the dynamics of house prices and mortgage debt when the economy is hit by a technology or monetary policy shock (Greenwald 2018, Castellanos, Millard, and Varadi 2025).

countries (e.g. in the UK the LTI limit is 4.5). Consequently, it bind for many prospective buyers, leading to strong effects even if it would have been imposed just in isolation.

C.3. Interactions between interest rate shocks and the tightness of credit limits

To understand how borrower based macroprudential limits interact with other shocks, we study the long and short run effects of our interest rate experiment under two credit regimes: *(a)* a *loose credit* regime in which LTV and LTI limits correspond to those prevailing in Ireland before the 2015 reform, and *(b)* a *tight credit* regime in which institutional limits are imposed.

Long run effects. Table A2 shows the long run effects of a 20 basis points permanent increase in the real interest rate before and after having imposed the institutional limits introduced in 2015. It indicates that the fall in the average house price and the homeownership rate is weaker when credit limits are more stringent, while rental rates continue to rise by similar magnitude. The later could be explained because the marginal landlord that drives the rental price response is likely to be unconstrained by those stringer limits, and hence reacts similarly to the real interest rate shock.

TABLE A2. A permanent increase in the real interest rate under two credit regimes

	Loose Credit (Pre-reform)	**Tight Credit (Post-reform)**
Δ Rent-to-Price	+10 bps	+10 bps
Δ%Average house price to income	-0.61 %	-0.4 %
Δ% Rent to Income	+1.44 %	+1.5 %
ΔHomeownership rate	-0.48 p.p	-0.3 p.p.

NOTE. This table shows the long-run effects of a 20 bps permanent increase in the real interest rate for two economies: *(i) loose credit* in which LTV and LTI limits correspond to the pre-reform economy, and *(ii) tight credit* in which institutional LTV and LTI limits are imposed and hence credit conditions are equivalent to those in the post-reform economy.

Short run effects. Figure A5 depicts the response of rental and house prices in deviations from steady state after the pre-reform (loose credit) and post-reform (tight credit) economies are shocked with a 100 basis points transitory shock, modeled as an AR(1) with 0.10 annual persistence. Similarly to the long run effects explained just above, the tighter credit conditions weaken the response of house prices to the shock. However, the two economies display a very similar response in rental prices. This again highlights that the marginal landlord is not effectively impacted by the credit tightness as rental prices need to temporarily rise by a similar amount in both economies.

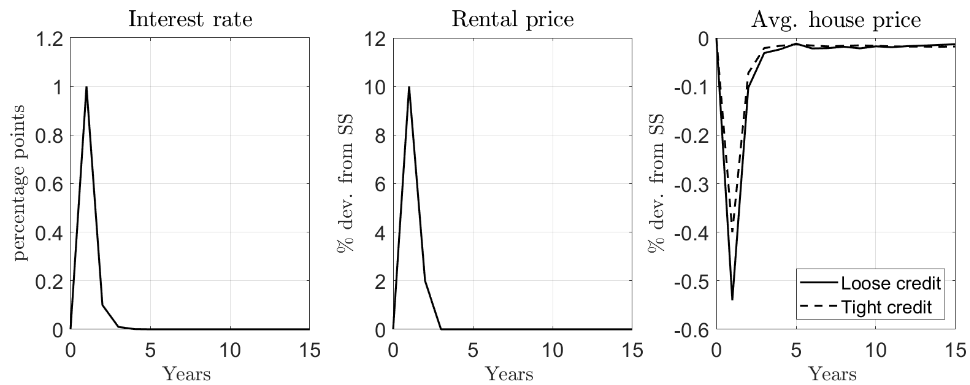

FIGURE A5. A temporary increase in the real interest rate under two credit regimes

NOTE. This figure shows the evolution of the interest rate, the rental price and the average house price after the pre-reform (solid line) and post-reform (dashed line) economies are shocked with a transitory increase in the real interest rate.

Implications. These two exercises highlight that macroprudential policies are effective in moderating house price cycles, consistent with the consensus in policy spheres. Hence, the model is still able to showcase some of the benefits associated to borrower based measures although it lacks its benefits associated to weakening aggregate demand externalities and reducing households' defaults. As a result, we do not investigate further what would be the optimal LTI/LTVs ratios that insure borrowers against the negative impacts of the shock.

Appendix D. A Model with Conversion Costs

The model in the main body of the paper takes costlesss conversion between housing types as its baseline assumption while Section 6.2 takes the opposite extreme by assuming housing types are totally different goods with different production functions and, therefore, no conversion between types. Here, we lay out a simple, micro-founded yet tractable way to introduce conversion costs into the model. We find that, even at relatively large assumed conversion costs, the transition dynamics are broadly similar to those we report in our main experiment.

D.1. Arbitrage and housing conversion with construction input

Suppose the existence of an arbitrageur who is able to buy, convert, and sell different housing types. We suppose an arbitrageur who is able to take advantage of the need to refurbish many houses at once and so, unlike a household, is able to repackage a measure of the housing stock into houses of different quality types. In our economy we have n housing types indexed by $i = 1, \ldots, n$, ordered by quality $\tilde{h}_1 > \tilde{h}_2 > \cdots > \tilde{h}_n > 0$, with the price of a whole type-i house is p_i^h. In the frictionless benchmark with perfectly convertible housing types, no-arbitrage equalizes price per unit of quality (we could think of this as price per square meter), i.e.

$$\text{(A6)} \qquad \frac{p_i^h}{\tilde{h}_i} = \frac{p_j^h}{\tilde{h}_j} \quad \text{for all } i, j \qquad \Longleftrightarrow \qquad p_i^h = \frac{\tilde{h}_i}{\tilde{h}_j} p_j^h.$$

We will now consider the possibility of conversion costs. We suppose that to convert a measure of housing x_{ij} from type-i to type-j, the arbitrageur must buy inputs from the construction sector. Specifically, let $\xi_{ij}(x_{ij})$ be the quantity of intermediate input required for the $i \to j$ conversions. The arbitrageur purchases these inputs at the prevailing cost of construction sector output p^h. Let $B_i \geq 0$ be the number of type-i houses *bought*; $S_i \geq 0$ be the number of type-i houses *sold*. Since the arbitrageur is merely repackaging the existing housing stock into different housing types, they face a feasibility constraint given by

$$\text{(A7)} \qquad g_i(S, B, x): \; \tilde{h}_i S_i - \tilde{h}_i B_i + \sum_{j \neq i} \tilde{h}_i x_{ij} - \sum_{j \neq i} \tilde{h}_j x_{ji} \;\leq\; 0, \qquad i = 1, \ldots, n.$$

This says that the quality-weighted number of houses of type-i sold is less than or equal to the number of houses of type-i bought less the number of type-i converted into other types plus the number of other types converted into type-i.

Problem. The arbitrageur is risk neutral and solves

$$\text{(A8)} \quad \max_{B,S,x\geq 0} \pi = \sum_{i=1}^{n} p_i^h(S_i - B_i) - p^h \sum_{i\neq j} \xi_{ij}(x_{ij}) \quad \text{s.t.} \quad g_i(S,B,x) \leq 0 \; (i = 1,\ldots,n).$$

Let $\mu_i \geq 0$ be the multipliers on (A7).

FOC for trades. The Lagrangian is $\mathcal{L} = \sum_i p_i^h(S_i - B_i) - p^h \sum_{i\neq j} \xi_{ij}(x_{ij}) - \sum_i \mu_i g_i$. FOCs with respect to S_i and B_i give

$$\text{(A9)} \qquad p_i^h - \mu_i \tilde{h}_i = 0 \;\; (\text{if } S_i > 0)$$

$$\text{(A10)} \qquad -p_i^h + \mu_i \tilde{h}_i = 0 \;\; (\text{if } B_i > 0) \;\Rightarrow\; \mu_i = \frac{p_i^h}{\tilde{h}_i} \quad \text{(whenever trades clear balances).}$$

Thus μ_i is the shadow value *per unit of housing* of type i.

FOC for conversion flows. For each $i \neq j$,

$$\text{(A11)} \qquad \partial_{x_{ij}}\mathcal{L} : -p^h \xi'_{ij}(x_{ij}) - \mu_i \tilde{h}_i + \mu_j \tilde{h}_i \begin{cases} = 0 & \text{if } x_{ij} > 0, \\ \leq 0 & \text{if } x_{ij} = 0. \end{cases}$$

Substitute (A9) to obtain the pairwise no-arbitrage condition :

$$\text{(A12)} \qquad p_i^h \geq \frac{\tilde{h}_i}{\tilde{h}_j} p_j^h - p^h \xi'_{ij}(x_{ij}), \qquad \text{with equality if } x_{ij} > 0.$$

Importantly, in spite of potential arbitrage gains (i.e.: when $\Delta_{ij} := \frac{\tilde{h}_i}{\tilde{h}_j} p_j^h - p_i^h \neq 0$), it is possible that the presence of conversion costs may mean there is a region where it is not worth the arbitrageur converting housing types. In this case, the stock of that housing type would essentially be fixed and we would need to search for the price that lead to demand equaling this fixed supply. For this reason, it is useful to restrict ourselves to cases where ξ is smooth and differentiable for all $x_{ij} > 0$ and that $\xi'_{ij}(0) = 0$. This has the spirit of an Inada condition and is present simply to avoid this corner solution. With this condition in hand, we can be sure that the arbitrageurs will always exploit price differences and ensure that a demanded change in composition of the housing stock will be met.

Example: Three Housing Types. With $n = 3$ the six directed links satisfy

$$p_1^h \geq \frac{\tilde{h}_1}{\tilde{h}_2} p_2^h - p^h \xi_{12}'(x_{12}), \qquad p_1^h \geq \frac{\tilde{h}_1}{\tilde{h}_3} p_3^h - p^h \xi_{13}'(x_{13}), \tag{A13}$$

$$p_2^h \geq \frac{\tilde{h}_2}{\tilde{h}_1} p_1^h - p^h \xi_{21}'(x_{21}), \qquad p_2^h \geq \frac{\tilde{h}_2}{\tilde{h}_3} p_3^h - p^h \xi_{23}'(x_{23}), \tag{A14}$$

$$p_3^h \geq \frac{\tilde{h}_3}{\tilde{h}_1} p_1^h - p^h \xi_{31}'(x_{31}), \qquad p_3^h \geq \frac{\tilde{h}_3}{\tilde{h}_2} p_2^h - p^h \xi_{32}'(x_{32}), \tag{A15}$$

with equality on active links. Along an active chain $i \to j \to k$, costs add with the appropriate size scaling via (A12).

Quadratic conversion costs. Suppose the input requirement itself is quadratic, $\xi_{ij}(x) = \frac{1}{2}kx^2$ with $k > 0$. Then conversion cost is $p^h \xi_{ij}(x)$ and the marginal cost entering (A12) is $p^h \xi_{ij}'(x) = p^h kx$. Because $\xi_{ij}'(0) = 0$, any strictly positive wedge Δ_{ij} induces the flow $x_{ij} = \Delta_{ij}/(p^h k)$ and satisfies

$$p_i^h = \frac{\tilde{h}_i}{\tilde{h}_j} p_j^h - p^h k x_{ij}. \tag{A16}$$

The above equations define house price gaps but are silent on the actual level. To close the model, we assume that the flow of new construction can be installed as any of the housing types. As such, the construction company will always sell its output as the most expensive housing type k. Perfect competition in the construction sector will then ensure that the per-unit price of type k will equal the price of construction sector output such that $p_n^h/\tilde{h}_n = p_h$. All other prices are then set relative to this anchor.

D.2. Robustness to Conversion Costs

To illustrate the limited impact conversion costs have in our particular setup, we take a value of conversion costs that is at what we consider to be the upper bound of plausible such that cost of converting big houses into smaller houses is half the cost of building a new small house from scratch. Figure A6 shows the result. The middle panel shows

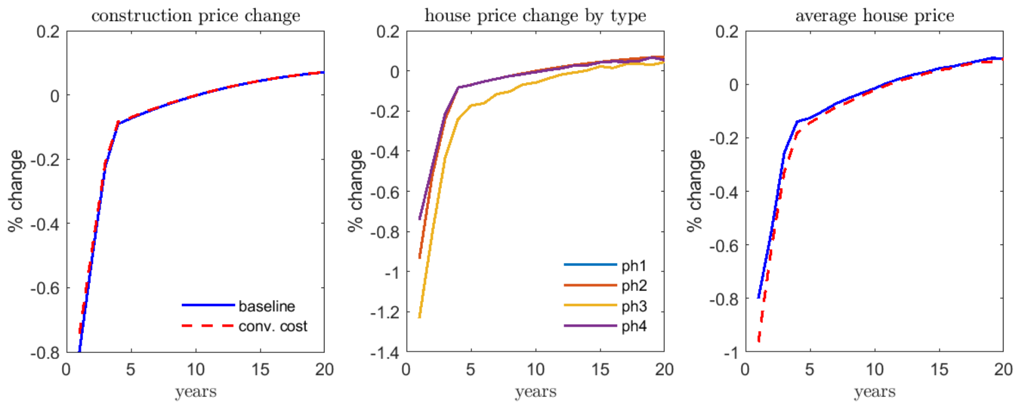

FIGURE A6. Baseline vs. Conversion Costs

NOTE: This figure compares the baseline transition with a version of the model with very high conversion costs.

the theoretical possibility of price dispersion between the houses that are relatively in demand vs those being moved away from. Despite this, the other two panels show that both the unit price is only very slightly higher (reflecting the fac that the construction sector must now support conversions) while the average house price drops by 0.1p.p. more on impact (reflecting the fact that the housing being moved away from drop my more than the unit price, pulling down the average).

The reason for this small impact is that, ultimately, the required number of conversions is low due to the fact that there is a flow of new housing being built and that these can be installed as any of the housing types - namely the ones subject to an increase in demand.