Dear Editor,

We are making use of the transfer option from the AER to AEJ:Macro following a recommendation from the AER editor, Ben Moll, as a response to our submission last year. In the meanwhile, we have revised the paper further, mostly as a response to the referee reports we received from the AER. While we believe the paper is much clearer and robust now, the changes to the paper have not affected its substance, with both the empirical application, the model and the key quantitative results being very similar to the original submission. Hence, we believed that the original reports, together with our responses to them, were still of relevance to the current version of the paper, and could aid your editorial decision.

In this cover letter, we first summarise the key changes that we have introduced to the paper since the AER submission, and then provide a detailed response to the issues raised by the referees and the steps we have taken to start addressing them. The point-by-point responses put this letter on the longer side but we hope it will be helpful as a guide to where the paper has been sharpened in response to comments. That being said, given that this was not a revision, there are some points which we have not covered. We also highlight these, together with a brief description of whether each point could be feasibly covered in a potential revision. We are of course available for any further clarification, should it be needed.

Many thanks for taking the time to review our manuscript. We hope you agree with us that it is a good fit for AEJ:Macro.

Gonzalo, Andrew and Juan

Key changes

1. **Model validation against the Irish empirical evidence**

   We have added Section 4.4. to the paper, in which we replicate our empirical results to those delivered by our model. Specifically, the empirical section is based on geographical heterogeneity across Irish counties; because our main results are based on a nationally representative model, we generate synthetic regions with different price-to-income ratios and hence different distance measures to be able to run a comparable regression. We find that the reactions of house prices and rents to the policy not only have the right signs in the model, but also have comparable magnitudes to those in the data, which we see as convincing evidence that the model is replicating the right elasticities and decision margins.

2. **Improving exposition of the paper and relationship to existing literature**

   Based on the advice of the referees, we have substantially improved the exposition of the paper.

   We have used the advice of Referee 3 to completely rewrite our Related literature section, and also to make the connection to previous literature clearer throughout our paper.

   We have used the advice of Referee 1 to more clearly establish the difference in mechanisms between our paper and the recent paper by Greenwald and Guren (2024), both in the intuition section (2.4) and in specific robustness sections (6.1. to 6.3.).

   We hope the main idea and contribution of the paper to the literature is much clearer as a result.

3. **Highlighting differences between short-run and long-run effects**

   We study both the effects of macroprudential policies when comparing two steady states and over the transition. In this new revised version, we use the results coming from both to more clearly described the short-, medium- and long-run effects of credit shocks, focusing on the differential reactions of rents and house prices at different horizons.

Responses to referees

Referee 1

On the empirical side, Gete and Reher (RFS 2018) already document that rents rise after a credit expansion using U.S. data. Appendix table B.1 in Greenwald and Guren (2024) also has a similar result (though not statistically significant). While there is value in confirming important findings in alternative settings, I think their results set the bar higher for a paper that does something similar. On the modeling side, while some of the modeling details are different from Greenwald and Guren (2024), my intuition is that the key mechanism is identical in both models - there is a cost of switching supply from rental to owner-occupied or vice versa. I conjecture that in that model, rents also go up after a contraction (they do not report rents, but Rent/Price goes up after a contraction). Despite this, I think it's a nice paper. I propose several ways to strengthen it.

The updated version of the paper now describes in much more detail why our mechanism is distinct to the one described in Greenwald and Guren (2024). While we will elaborate on this point as part of our response to some of the specific issues that are raised, in pages 5-6 we now write:

"In our model, heterogeneity in age, income and wealth, together with financing constraints and the lumpiness of housing, provide a micro-foundation for the upward-sloping rental supply curve, driven by the marginal landlord requiring additional compensation to buy more housing. This approach contrasts with upward-sloping rental supply curves generated by exogenous preference distributions (Greenwald and Guren 2025) or intermediary rental supply companies (Dias and Duarte 2022, Rotberg and Steinberg 2024). Moving away from reduced form mechanisms allows us to study the impact of policy changes on this supply curve differentiating between shocks that mostly impact rental demand (borrower-based macroprudential measures) and those that impact both demand and supply (real interest rate shocks). Our setup also allows us to separately study the impact of shocks on house prices and rents, contrasting them with the empirical evidence on each dimension. This distinction matters because the welfare consequences of a given change in the rent-to-price ratio depend on the underlying source of that

movement – whether it reflects shifts in rents, prices, or both. In our model, credit shocks generate an increase in rents, even with a sizable house price drop, and are associated with a transfer of welfare from young renters towards older landlords"

Additionally, in Section 6.1. we show that a model with no segmentation between owner-occupied and rental properties still generates increasing rents as a result of credit shocks. While quantitatively, the amount of segmentation matters, particularly for the response of house prices (see Section 6.3), this observation highlights that our mechanism is distinct to the one described in their paper.

1) Empirics: I think you could add value by showing that your key finding – rents going up (down) after a credit tightening (expansion) is a robust finding that happens in various settings. You could do this by following Greenwald and Guren's strategy. They use three different IVs for credit shocks that have been previously used in the literature: Loutskina and Strahan (JFE 2015), Di Maggio and Kermani (RFS 2016), Mian and Sufi (RFS 2022) to check how a credit expansion affects house prices and home ownership. You could do the same for rents. Perhaps you could identify even more similar IVs. I think there would be a lot of value in showing that this effect exists and how strong it is across settings. A separate thought I had: How quickly do rents rise after a credit contraction? If it's a slow reaction, maybe it has more to do with aggregate supply growth slowing due to lower house prices rather than the channel you have in mind?

We have chosen to keep the focus on the Irish case and on the macroprudential policy that represents our main policy experiment, and have worked on strenghtening the connection between model and empirics. We connect our main experiment to our real interest rate shock exercise and describe how similar mechanisms are at play. This is very much aligned with recent empirical literature showing interest rate rises also increase rents, which we see as additional evidence for the robustness of the mechanism, despite the different nature of the shock.

"Gete and Reher (2016), in a two-period model, provide theoretical support to the idea that restricting mortgage credit increases the demand for rental housing, which is consistent with the empirical evidence in Gete and Reher (2018) that the contraction of mortgage supply after the Great Recession increased housing rents. Recently, Dias and Duarte (2019) and Abramson, De Llanos, and Han (2025) have also found empirically that contractionary monetary policy shocks increase rental prices." (page 6)

We have also expanded the discussion of short vs. long-run effects (section 4.3.2) and show (e.g. in Appendix C.1.) how conceptually our channel is also present in the short run and, if anything, is stronger because of the relatively lower elasticity of landlords, who take time to buy new housing. Although in the short-run house prices drop by almost 1 percent, our rent effects are also present in a version of the model in which there is no segmentation nor quality heterogeneity and hence prices do not move (Section 6.1.).

2) Model: My overarching comment is that the model is more complicated than necessary and there would be value in simplifying it. I think you would get similar results but it would be more transparent.

We have sharpened the model description and provide separate analysis of the effects turning certain model elements off or adopting different assumptions. We address specific concerns below.

2A) I don't understand the purpose of different housing qualities, especially if their prices all perfectly scale. For example, in Landvoigt, Piazzesi, Schneider (AER 2015) there are different qualities, but their relative prices change with the local supply/demand for each. I guess this is to have renters and owners? But I believe you could do this in a simpler way by having a single housing type and a utility preference for owning. In this case, people with high expected income growth and low current wealth will choose to rent because they cannot meet the LTV and LTI requirements. In fact, you already have the utility boost theta. Furthermore, since there are just two property types, is there any difference at all between your model and one where there is simply a preference for ownership in the utility function and one type of housing? I conjecture that not really.

In Section 6.1. we report results for a model with a single housing quality. We conclude that our results for rents are qualitatively similar in that framework. However, heterogeneous qualities are necessary to get house prices to move as a result of a credit shock in our framework, because otherwise aggregate housing stock would be constant (more on this in point 3B below). To show that our rising rents prediction survives the possibility of falling house prices, we have chosen to keep the heterogeneous housing types for our main results.

We have also clarified how housing tenure and housing qualities operate in our framework and enter the utility function of households (e.g. with Equation 2).

2B) I find the way you model depreciation cumbersome and I doubt doing this is

important for your result so you should simplify. Why not just have a single type of depreciation $delta_h$, that is get rid of $1/J$?

Footnote 5 clarifies this issue: "This assumption also captures the fact that frequently elderly households occupy housing that, because of its characteristics and location, may have become less desirable over time. Hence, it is also a simple way of taking into account that there is additional housing production in the economy e.g. because of aggregate trends of urbanisation and geographical reallocation, without assuming that housing depreciation δ_h, which in our model is paid directly by homeowners every period, is unreasonably high."

2C) Do you need the housing transaction costs? I agree that housing is illiquid and transaction costs are realistic, but if they are not central to your model's results, I think you may be better off getting rid of them to make it easier for the reader to understand what is going on in the model. Same for housing quality types.

We believe we do need them. The ease with which landlords can enter/exit the market is a key determinant of the rental supply response, especially in the short-run. To capture this, we have kept housing transaction costs, but we have extensively clarified the discussion on them and other features of the housing good.

2D) The effect on prices is surprisingly small. For example, Greenwald and Guren have a much larger effect of credit conditions on prices in a model that should have the same key features. I don't fully understand what is going on. I think one thing that might be driving this is the assumptions on the functional form of construction firms, combined with there being no risk, and a constant housing supply is constant. This is what allows you to have an analytic solution for the house price as a function of just quantity of housing and housing sector productivity Ah. I suspect this drives some of the results. For example if the return on housing must be equal to r, and if pH is fixed, then rent must rise when r rises. But in alternative specifications, you might get rent staying constant and pH falling instead. For example, imagine a very simple model where housing expenditures are always a fraction alphaH of income, income is constant, and there is no depreciation. Then rent is constant and the house price is simply PV of future rents, which falls when rates rise. Furthermore, equation 20 would be different if the housing stock was changing. In that case $d_h * H$ would be replaced by $d_H * H + (Hnext - H)$ so if interest rates are rising and investment in housing becomes less attractive, the housing

supply will fall, this quantity would be smaller than $d_H * H$, and the price would fall as well.

We have now clarified the section that describes construction (2.2.2) and removed many of the points that could have led to confusion.

We now focus much more on the comparison between short-run and long-run effects, and show that the impact on house prices in the short-run is sizeable (-0.9%, Figure 6), while the long-run effect is very small. We now explain that, while in the long-run construction adjusts to keep the housing stock relatively stable, in the short-run there might be larger mismatches between supply and demand at a given price. Thus, despite the fact that we do not allow for time-to-build in our framework, we capture some of the key features of the housing market.

In point 2Ei below we clarify with respect to the separate reactions of house prices and rents (which are a key part of the contribution of the paper) and in point 3B we clarify with respect to the aggregate housing stock.

2E) The interest rate rising experiment.

2E.i) When interest rates rise, why does rent rise instead of prices absorbing everything? For example, consider a model where rates have no effect on supply / demand and therefore on rent. In this model, the price is the present discounted value of rents, and it would fall, absorbing 100% of the rate rise. Consider a more complicated model where rent is determined by supply/demand. Supply is unchanged in short run. Demand will rise or fall depending on whether income or substitution effects dominate. Thus, the change in rent will fully depend on whether IES is above or below one (I solved a simple 2-period model and confirmed this).

We have revamped the description of our model intuition, in Section 2.4., to clarify how housing and rental supply curves work in the model, for both the macroprudential and the interest rate experiments.

The bottom line is that, in this model, because housing tenure is endogenous both for renters and for landlords, a natural separation of the equilibrium of the housing and the rental markets arises because of the heterogeneity in stochastic discount factors and the presence of borrowing constraints. Take the case of an increase in the real interest rate. In a simpler model without heterogeneity, it would translate one-to-one to the rent-to-price ratio, and whether changes come from rents or prices would not necessarily be

determined. In our model, the case in which prices fall and rents do not move might not be an equilibrium because there might be excess supply in the rental market (too many landlords are willing to buy at that price) or excess demand in the rental market (not enough landlords are willing to buy at that price because they are constrained or because it is very costly for them to forego current consumption). Both housing and rental markets clear simultaneously in equilibrium.

This key point is also highlighted in the comparison with Greenwald and Guren (2024) in the related literature section (pages 5-6, also cited above).

2E.ii) What causes the rise in rates? The effect on endogenous quantities may differ depending on the reason for the rise. Is this a sufficiently interesting experiment? I think for an open economy like Ireland its fine, for the US maybe not

We agree that this is a relevant question, but we have chosen to focus on an exogenous rise in the real interest rate (i.e., coming from world markets) for simplicity. Our Ireland application reassures us that this is not an unrealistic assumption, but we agree that this might not be the case for the US.

2F) Calibration 2F.i) Living with certainty until age 95 seems unrealistic. You should have a shorter lifespan or a probability of death.

We have kept the age structure of the model as it was for simplicity, but this change would be relatively easy to incorporate.

2F.ii) You are targeting a rental supply elasticity of 1.4, but you get 3.5. I don't understand why you miss by so much. I presume that the higher the cost of switching from owner-occupied to rental, the lower the elasticity will be. I didn't see this parameter discussed anywhere. So all you need to do is raise that cost. Or maybe I am missing something.

We now describe the comparison between 1.4 and 3.5, which is part of our model validation and is not a targeted moment, in much more detail in Section 4.4.

Specifically, we write:

"However, this comparison is imperfect for several reasons. First, it relies on US data, where rental contracts, the construction sector and the availability of land may be different to Ireland. Second, it is estimated based on a policy (property taxes) which is different to our focus (credit shocks). Third, and most importantly, they are long-run elasticities, while in our model we are mostly concerned with the effects of a credit

tightening over the first few years after the reform"

That is why we have now switched to an alternative model validation which is closer to our main model experiment, namely, replicating the Irish empirical evidence in our model (the remainder of Section 4.4.).

Unlike Greenwald and Guren (2024), our model does not feature an explicit cost of switching between owner-occupied and rental units that can be adjusted to match a potential calibration target. The cost of additional rental units is related to the availability of willing endogenous landlords to provide them, which is a complex endogenous object in our framework.

However, there are features of the model that could be adjusted to match the 1.4 estimate. These include the initial distribution of wealth, preference parameters such as the discount rate, considering heterogeneity in rates of return, the distributions of preference shocks, etc. We have chosen not to do this for the reasons described above (it is not the exact counterpart to our experiment) and because we are reassured by our new model validation, but we conjecture that matching the estimate should be feasible with reasonable modifications to our model, at the risk of not matching the Irish empirical counterpart anymore.

3) Exposition

3A) Through at least the end of section 2.3, it is unclear to me if you have aggregate shocks in the model or not. I think you do not, but you should clarify. If you do not have aggregate shocks, then Ac and Ah are constants and you do not need both of them, or maybe you need neither, in either case you should specify that they are constants.

We now clarify in Sections 2.2.1 and 2.2.2. that A_c and A_h are constants. We need at least one of them to be able to replicate the empirical level of the house price to income ratios.

3B) It appears that the housing stock in the model is constant. In this case, you are confusing the reader when you talk about housing output being used for new construction. Related to this, I don't know how exactly you do the MIT shock transition. However, if you are using equation 20 then you have a mistake. This is because equation 20 is derived by assuming that construction demand is just equal to depreciated capital, which is not true when the housing stock is changing.

Our previous discussion of the construction sector was very unclear - we have com-

pletely rewritten that section. We now show clearly that the aggregate housing stock is not constant in our economy. We write:

" The aggregate housing stock, i.e. the aggregate amount of the housing good, can be measured as the quality-weighted sum of all housing units in the economy. If we let $H^{sh}_{n,t}$ denote the share of houses of quality n,

$$H_t = \sum_{n=1}^{N} H^{sh}_{n,t} \tilde{h}_n \ . \tag{1}$$

"

As it becomes apparent later on, changes in policies induce changes in the distribution of housing qualities, which change the housing stock. This is the mechanism through which house price changes. We have dropped the old equation 20 that led to confusion.

3C) There is no difference in housing size, just quality. This is explained in the housing choices section on pg 10, however it would have been helpful to at least mention earlier (when you introduce quality) that there is no size. I was confused and trying to figure out how housing works between pages 8 and 10. You talking about price per square foot on pg 8 made this extra confusing.

Our new Section 2.2.2. is hopefully much clearer. Although we refer to $\tilde{h}$ as housing quality in a broad sense, nothing precludes an interpretation in which it only represents size, as long as Equation (14) holds (a definition of the aggregate housing stock).

3D) I am confused by the utility function for housing. You write $f(h, h_h at)$ where h is houses owned, and $h_h at$ is the house you live in, but then when you define $f(h, h_h at)$ in equation 10, it seems to not depend on h, only $h_h at$ – this makes sense utility should not depend on the kinds of investment assets you own. So if I am understanding this correctly, you should just write $f(h_h at)$.

We have rewritten our model description and clarified this issue with the new Equation (2).

3E) You introduce depreciation on pg 8 and then basically repeat yourself on pg 10, I suggest deleting this on pg 10

We have rewritten our model description. We still kept a mention of depreciation in the household's problem, to clarify that homeowners pay for the costs of depreciation every period, and then describe in Section 2.2.2. how the housing services that provide

this reduction in depreciation are produced.

3F) footnote 2 mentions that the LTV and LTI constraints may potentially be different for first time buyers vs investors. I don't know about Ireland, but in the US, the mortgage interest investors pay is typically much higher than owner-occupiers, so if you allow LTV and LTI to differ, kappa should differ to. Having said that, I am encouraging you to simplify the model to the simplest you can get that still delivers the result, so I am by no means insisting you add all these differences to the baseline model. Perhaps you can have an appendix model where everything is more realistic (differing kappa, LTV, LTI; transaction costs; different house qualities . . .)

For simplicity, we have chosen to keep the mortgage interest rates constant for all individuals.

3G) The intuition in section 2.4 is wrong. You write that an increase in the rent/price ratio shifts landlords' curve up. No! It shifts landlords up the curve (that's what a supply curve is, you supply more when the price is higher). It shifts the households' curve not because it changes prices, but because for the same price, renting becomes more attractive.

You are right and the new Section 2.4. does not contain this error. Thanks for pointing this out.

3H) You make a point of your model having endogenous landlords, etc. However, this is less of an innovation than you let on. Your model isn't quite GE or PE because the housing market clears but the interest rate is exogenous. Since there is no risk, landlords must earn r on their real estate investments. This means landlords are indifferent about investing in real estate or in liquid assets. So it doesn't matter whether you model endogenous landlords (as you do) or if you had an alternative model where households only have access to a liquid asset and all landlords are outside of the model, as long as the rate of return on housing is r, I conjecture you will get exactly the same result.

We have hopefully clarified with the rewriting of our paper and with the new Section 2.4. that this is not the case. It is the heterogeneous stochastic discount factor of landlords and their different willingness to provide rental accommodation that generate our upward-sloping supply curve for rental housing at a given house price. Therefore, we would not get this result if landlords were outside of the model.

We make this explicit with new Figure 2. We write:

"The negative credit shock shifts the rental demand curve outwards, as shown by the

blue dashed line in Panel B of Figure 1: some potential buyers no longer qualify for a mortgage after the reform, so they cannot buy a house anymore and are pushed back into rental market. This increase in rental demand interacts with the upward sloping supply curve: the negative credit shock not only leads to a reduction in the homeownership rate, but also to an increase in the rental prices these new renters must pay. Figure 2 shows where the movement along the supply curve comes from: the new additional rental units must be provided by new landlords with higher rent supply thresholds than the existing ones"

3I) In intro, I don't understand their explanation ii). If everyone moves to smaller properties, that will reduce rents and prices if there are no other frictions.

We have clarified this further. However, we agree that explanation (ii) in the introduction would push down both house prices and rents as people move to smaller properties. We chose to focus on the effect on house prices because that is what we observe in our empirical evidence and what the model replicates. In our simple model, there is no heterogeneity in rental properties and therefore rents cannot drop from this channel; we chose not to include this according to our empirical evidence, in which rents rise as a result of both the macroprudential and the interest rate shock. Adding heterogeneity in rental properties would be possible, but computationally costly.

Referee 2

1. The objective of this paper is to analyse the distributional implications of credit shocks on housing and rental markets. Policies on maximum loan-to-value (LTV) and loan-to-income (LTI) should affect differentially wealthier households (LTV) or higher income households (LTI), and age correlates with wealth and income. Nevertheless, the empirical part on the Irish data is not exploiting e.g. the wealth differentials stemming from LTV incidence. Therefore, the authors should exploit further the data to understand how rental and house prices change depending on e.g. wealth or age (the distance measure used is not related to e.g. wealth or age). In addition, the action of buying a house may imply changing the county where the household lives renting an apartment (and thus buying in a different county), this is not addressed either. Improving this part is essential given the objective of this paper.

We agree that our model takes a number of simplifying assumptions, including the fact that Ireland is a single housing market and we do not model geography explicitly.

However, in our new replication of the empirical evidence within the model (Section 4.4), we generate a regression identical to the one we estimate in the data by solving the model for a plurality of counties with different house price to income ratios, hence moving one step further to adding geographical heterogeneity. We see endogenous moving across counties as a natural next step, but beyond the scope of our current paper.

With respect to the empirical part, there is no information about household wealth in neither the original credit registry used in Acharya et al. (2020) nor in the mortgage data that we have access to. Therefore, we cannot exploit differences in wealth. We do have information about age, but we chose to keep the empirical part as simple as possible to be comparable with what had already been done in the Acharya et al. (2020) framework.

However, to the extent that households with different income and wealth will endogenously choose different LTIs and LTVs, which would then aggregate to different county-level distance measures, our empirical estimates do incorporate some of the heterogeneity we observe in the data. In the model, both income, age, and wealth are explicit, and we show (e.g. Figure 4) how some of these untargeted moments compare in the model and in the data.

2. Less essential, but nevertheless, crucial for household welfare is to understand whether there is any benefit of these policy (credit) changes. That is, these policy changes

happen because there are intended positive consequences, which could be higher for households with less income and wealth (younger). There is an influential literature by e.g. Mian and Sufi that argue that subprime borrowing was crucial for the 2007-09 crisis (externalities here could be fire sales, credit crunches). In Europe (Ireland), effects could be even higher than USA for households welfare given that mortgage loans are with full recourse (externalities here could be not only via fire sales but also massive consumption reduction). Therefore, the paper could analyse both the benefits and costs of policies on credit, as this will have for sure importante effects for household welfare, especially distributional effects.

We agree that this is an important point that we do not tackle in our current paper, and we see it as a priority for follow up work, as it requires a substantially different modelling framework and, in particular, the addition of some form of aggregate risk.

In the introduction we write: "Nonetheless, we cannot measure the benefits of those policies given that we do not model the cyclical buildup of risk in the financial sector and, as a result, we are silent on optimal LTV or LTI ratios."

In the conclusion we write: "A broader framework with a rich financial sector could jointly evaluate the benefits of macroprudential policy coming from increased financial stability with the distributional costs imposed on households."

3. Much less important, and not essential, but interesting for macro is that once credit is changed and hence house prices are changed, this implies different implications for firms' access to credit. Smaller firms will be differently affected, and lower income households may work substantially more at smaller firms, with important distributional implications.

This is also a very interesting point, which which we agree, but which we see as beyond the scope of our current framework. (Add to conclusion??)

Referee 3

We would like to extend our thanks to Referee 3 for his very extensive report, which has substantially improved our paper and also our understanding of the previous literature and how our paper compares to it. In the interest of space, we do not copy it here. We just copy the "Recommendations" which more closely relate to our paper.

1. Rental market supply elasticity.

Recommendation: I would encourage the authors to familiarize themselves with the early developments in the literature to ensure they have a thorough understanding of how rental prices respond to shifts in the supply and demand of rental units. The references in the paper only address very stylized cases (e.g., Guren and Greenwald, 2024), and their focus is solely on whether rental units are perfect substitutes for owner-occupied units. This is a direct response to the paper by Kaplan, Mitman, and Violante (2020), which presents a narrow and incomplete theory of house prices. I argue that their work is a distraction, as other research has demonstrated that house prices can be aligned with data in response to changes in credit conditions. These movements can occur with or without segmentation (see below for a pencil-and-paper specifcation). For example, Garriga and Hedlund (AER 2020, 2019) have shown this to be the case.

We have reorganised our Related literature section to better describe previous contributions to the macro-housing literature that are of relevance to our framework, particularly those that have dealt with rental supply elasticities. We have also added many of the references that were pointed out in the report. We think the description is now much more complete and are grateful for the clarification.

We have also used these insights to re-focus our introduction and the description of our key contribution to the literature.

2. House prices and regime changes/duration of credit conditions

Recommendation: The take away from this exercise is that in this setting the critical factors such as where are current rates before the change in policy (low rates or high rates, as this determines the distribution of credit). I think the authors could do more connecting the levels, and the persistence of the policy changes with the work already done in the literature

Our revised version of the paper is much clearer about both short-run and long-run effects of the policy reform, but we have kept the key structure of the reform unchanged

(a permanent unexpected change to LTV/LTI ratios or to real interest rates based on the same pre-existing steady state levels).

We very much agree with the referee on the importance of where rates and financial conditions have been in the years before any given reform, and also particularly given the specific circumstances in Ireland around the time of the Global Financial Crisis. We have not explored this avenue further to keep the paper contained and to have a clearer message. However, our framework allows to study the path dependence of interest rates and more complex policy reforms (including transitory and persistent ones).

In terms of the connection with previous literature, in the related literature section (page 6) we write:

"Garriga, Manuelli, and Peralta-Alva (2019) show that, even in a representative agent framework, a reduction in the mortgage rate may increase house prices and reduce rents in the short run. Lower borrowing rates induce both a positive income effect and a housing investment boom. If the latter is strong enough, the relative price of housing services (rents) might decrease in the short run. Our heterogeneous agent framework generates a complementary channel via the portfolio decisions of the owners of rental properties that leads to rents moving in the same direction as mortgage rates, even in the short run"

We also clarify in our results section that "We study the effects of the macroprudential reform under the assumption that it is a permanent change."

3. Endogenous house prices and rents in response to a change in credit conditions

Recommendation: The authors have the empirical discipline to conduct the experiment in Ireland and calibrate the model's response. However, it is important to be very clear and transparent about the indeterminate on house prices and rents from the basic theory.

We have replicated our empirical estimates for Ireland within the model and find that they are very close (Section 4.4.) without explicitly targeting them. This reassures us that the model is capturing the right margins, both for rental and housing demand and supply.

In the intuition Section 2.4. we have clarified how our model generates movements in rents and house prices separately, and what disciplines the supply and demand curves for both housing and rental markets, making this issue hopefully clearer.

In the related literature, we clarify how our model can deliver endogenous reactions

of both rental prices and house prices, which is not the case of more basic frameworks, as you highlight:

"In our model, heterogeneity in age, income and wealth, together with financing constraints and the lumpiness of housing, provide a micro-foundation for the upward-sloping rental supply curve, driven by the marginal landlord requiring additional compensation to buy more housing. This approach contrasts with upward-sloping rental supply curves generated by exogenous preference distributions (Greenwald and Guren, 2025) or intermediary rental supply companies (Dias and Duarte, 2022; Rotberg and Steinberg, 2024). Moving away from reduced form mechanisms allows us to study the impact of policy changes on this supply curve differentiating between shocks that mostly impact rental demand (borrower-based macroprudential measures) and those that impact both demand and supply (real interest rate shocks). Our setup also allows us to separately study the impact of shocks on house prices and rents, contrasting them with the empirical evidence on each dimension. This distinction matters because the welfare consequences of a given change in the rent-to-price ratio depend on the underlying source of that movement – whether it reflects shifts in rents, prices, or both. In our model, credit shocks generate an increase in rents, even with a sizable house price drop, and are associated with a transfer of welfare from young renters towards older landlords. "

4. Minor comments

4.1. Limited consideration of supply-side factors: The paper primarily focuses on the demand-side effects of credit shocks on housing and rental markets. While it acknowledges the role of housing supply, it assumes a fixed amount of buildable land (housing permits) each period. This simplification may not fully capture the complexities of housing supply responses to changes in demand and credit conditions. In reality, factors such as zoning regulations, construction costs, and developer incentives can significantly influence housing supply elasticity. A more comprehensive analysis would consider the dynamic interplay between supply and demand forces in response to credit shocks.

We agree with this comment. In footnote 10, we write "In reality, zoning regulations, construction costs and other determinants of housing supply may react endogenously to changes in demand and credit conditions, but they are likely to adjust relatively slowly. We leave this channel for future work."

Assumption of a single, national housing market: The model assumes a single, national

housing market, which ignores regional variations in housing market conditions, demand, and supply dynamics. This simplification may not accurately reflect the localized nature of housing markets, where factors such as local economic conditions, amenities, and population growth can drive significant price and rent differentials across regions. A more nuanced analysis would consider the spatial heterogeneity of housing markets and the potential for spillover effects across regions in response to credit shocks.

We also agree with this important simplification. As a small step in that direction, our replication of the empirical evidence within the model now allows for regional variation as well (i.e. we estimate the model for several submarkets with different price-to-inocme ratios). However, we see a broader model with more detailed regional variation and endogenous moving as an important next step. We write in the conclusion:

"Another interesting extension would consider spatial heterogeneity across different urban and rural areas, and the interaction of a credit shock in a context of accelerated urbanization and increased house price and rent inequalities"

Simplification of mortgage market features: The model makes some simplifying assumptions regarding mortgage market features, such as treating mortgages as negative asset holdings and imposing a fixed minimum amortization payment. These simplifications may not fully capture the complexities of real-world mortgage contracts, which often involve various features like prepayment penalties, adjustable interest rates, and the possibility of default. A more realistic model would incorporate these features and their potential interactions with credit shocks. Other papers have address the role of long-term contracts that separate the stock of credit from its flow.

We also agree with these limitations. We have tried to make our assumptions as clear as possible in the new model description (pages 9-10).

Absence of financial sector dynamics: The model doesn't explicitly model the financial sector and its interactions with housing and rental markets. One the obvious one is mortgage default in response to a negative credit supply shock that reduces prices. This omission limits the ability to assess the potential feedback loops between credit supply, housing demand, and financial stability. For example, a tightening of credit conditions could affect the availability of mortgage financing, leading to further reductions in housing demand and price declines. A more comprehensive model would incorporate financial sector dynamics and their role in amplifying or mitigating the effects of credit shocks.

We also agree with this limitation, which we highlight in the introduction, related literature and when we discuss our main results.

In the introduction we write: "Nonetheless, we cannot measure the benefits of those policies given that we do not model the cyclical buildup of risk in the financial sector and, as a result, we are silent on optimal LTV or LTI ratios."

In the conclusion we write: "A broader framework with a rich financial sector could jointly evaluate the benefits of macroprudential policy coming from increased financial stability with the distributional costs imposed on households."

Overall, while the paper provides valuable insights into the impacts of credit shocks on housing and rental markets, it's important to recognize the limitations of the analysis and the findings. Future research could address these limitations by incorporating more realistic features of housing supply, regional heterogeneity, long-term effects, mortgage market complexities, and financial sector dynamics.

We agree and acknowledge many of these limitations, suggesting future work in these directions in our conclusion.

References

Dias, Daniel A, and Joao B Duarte. 2022. "Monetary Policy and Homeownership: Empirical Evidence, Theory, and Policy Implications."Technical report, Working Paper, Board of Governors of the Federal Reserve System (US).

Greenwald, Daniel L., and Adam Guren. 2025. "Do Credit Conditions Move House Prices?" *American Economic Review* 115 (10): 3559–96.

Rotberg, Shahar, and Joseph B Steinberg. 2024. "Mortgage interest deductions? Not a bad idea after all." *Journal of Monetary Economics* 144.