Referee Report for “The Aggregate and Distributional Implications of Credit Shocks on Housing and Rental Markets”

1 Summary

This paper studies the equilibrium effects of contractionary credit shocks on housing and rental markets by building a life-cycle model with households that are heterogeneous in age, income, and wealth, and who make endogenous housing tenure choices (rent, own, or become landlords). House prices and rents are determined in equilibrium. Relative to many papers in the macro-housing literature, risk-neutral/absentee landlords are replaced by household-landlords inside the model with their own wealth, income, and financial constraints. Combined with frictions arising from the lumpiness and illiquidity of housing investment, this generates an upward-sloping rental supply curve with empirically plausible elasticities—that are an order of magnitude smaller than in other models,—which is the engine of the paper's main results.

The paper's central finding is that policies designed to cool housing markets, such as limits on loan-to-value and loan-to-income ratios and increases in the real interest rate, succeed in reducing house prices but simultaneously increase rents in equilibrium. When stricter mortgage loan caps prevent marginal buyers from purchasing homes, they remain in the rental market, increasing rental demand. Meeting this demand requires recruiting new landlords who are less wealthy than existing ones, and who therefore require higher rents as compensation for making a large, lumpy, illiquid investment. The result is a persistent increase in rents even if the decline in house prices is transitory, a decline in the homeownership rate, and a wealth and welfare redistribution from young, low-income renters to older, wealthier landlords.

The paper validates its mechanism using Ireland's 2015 macroprudential reforms, finding that counties where the reform was most binding simultaneously experienced slower house price growth and faster rent growth–the opposite of their normal positive co-movement. The model, calibrated to the Irish economy, closely replicates these empirical patterns. The welfare cost of the reform is estimated at 1.5% of lifetime consumption for newly born households. A second experiment studies the effects of a rise in the real interest rate, finding qualitatively similar results: rents rise, house prices fall, and homeownership declines. In both experiments, the authors analyze transitions to study the persistence of the effects.

2 Comments

1. **Marginal Contribution**
   I enjoyed reading the paper. It is competently executed and the empirical validation of the model t Ireland is a strength. The emphasis on generating the right rental supply elasticity in a structural model is a good pitch and a nice contribution. That said, my overall impression is that the paper's novelty does not rise to the level of a JPE publication. There has been an extra-ordinarily rich macro-housing literature over the past 15 years that has studied the aggregate and distributional impacts of macro-prudential housing market policies. A few prominent examples are Landvoigt, Piazzesi, and Schneider (2015 AER), Favilukis, Ludvigson, and Van Nieuwerburgh (2017 JPE), Kaplan, Mitman, and Violante (2020 JPE), Garriga and Hedlund (2020 AER), Favilukis, Mabille,and Van Nieuwerburgh (2023 RES), and Greenwald and Guren (2025 AER). Several of these papers clear both ownership and rental markets and think about the role of landlords that are internal to the model. I am not sure how much more appetite there is for another paper on macro-prudential policies in the housing market. The paper explains its contribution well, summarized best by the sentence "we provide a micro-foundation for the exogenous preference distributions which Greenwald and Guren (2025) use to generate realistic reactions of rent-to-price ratios and homeownership to credit shocks," but that contribution may be a better fit for a journal like *JPE: Macro*.

2. **Mortgage Lock**
   The results on the impact of increases in the real rate are potentially more interesting (than the LTV/LTI constraints) in light of the large real rate increases observed in 2022-24. However, the puzzling thing is that house prices did not fall in this period and neither did homeownership rates, at least in the United States. A growing literature on mortgage lock, which features owners and renters, is trying to explain why (e.g., Fonseca, Liu, and Mabille, 2026). This paper does not have the features required to speak to this debate: a distribution of homeowners with different long-term mortgages rates originated at different dates. This limits its ability to speak to the current period.

3. **Constant Interest Rates and No Aggregate Risk**
   The paper has no aggregate risk and interest rates are exogenous and constant, limiting the ability of the model to speak to housing booms and busts, and their impact on the real economy. In particular, the core mechanism in the paper is wealthy households' willingness to become landlords, which depends on the returns to renting out houses relative to investing in risk-free bonds. In reality, the returns to investing in rental housing are risky. In a model with aggregate risk, rental income yields and capital gain yields would fluctuate and there would be a risk premium on rental investment, driven by the covariance of those returns with the intertemporal marginal rate of substitution of the marginal landlord. If times in which real rates go up or in which macro-prudential policies become binding are good vs. bad states of the world, that could affect the magnitude of the risk premium. Either way, consideration of aggregate risk would potentially have meaningful impact on the core mechanism and main results.

4. **The Missing Financial Sector**
   A main weakness of the paper is the absence of a financial sector. The raison d'être for macro-prudential policy is financial stability, to avoid the deadweight losses associated with a banking sector blowup. This means the welfare cost of 1.5% of lifetime consumption is estimated in a vacuum, without consideration of the financial stability benefits the policy is designed to deliver. This is not a minor omission since Ireland had a major recession following its financial sector collapse in 2008, directly attributable to excessive mortgage leverage of the kind that credit contractions are designed to prevent. The welfare gain from avoiding that crisis (which saw high unemployment, large real estate wealth destruction, and massive bailouts by the fiscal authority, etc.) almost certainly dwarfs the 1.5% welfare cost number highlighted in the paper. The authors acknowledge this limitation openly, but do not attempt any quantification of the financial stability benefits since it would significantly complicate the structural model. One would need to think about mortgage default and bank bailouts, and keep track of banking sector wealth as an additional continuous state variable. Some papers in the literature have gone down this path, e.g. Elenev, Landvoigt, and Van Nieuwerburgh (2016 JME, 2021 Econometrica). Seriously incorporating a financial sector would substantially elevate the contribution of the paper. Absent that, some discussion of what the literature suggests the welfare costs of financial crises is, and of the impact of macroprudential policies on financial crisis probabilities, would be helpful.

5. **Minor Comments**

   (a) The model assumes all mortgages are adjustable rate, meaning interest rate changes transmit immediately to all existing mortgagors. This is calibrated appropriately to Ireland in 2015, but it significantly limits the paper's interest rate experiment as a guide to monetary policy transmission in countries such as U.S. and several continental European countries where fixed rate mortgages dominate. The paper should be clearer about this scope limitation, particularly given its stated ambition to speak to the 2022–23 monetary policy tightening cycle, which operated across economies with very different mortgage market structures.

   (b) It is not clear what time period the calibration targets in Table 2 refer to. Are the pre-reform averages?

   (c) I would like to see an extra column in Table 4 with the numbers from the data.

   (d) The 1.5% CEV figure is reported for newly born households under the baseline calibration, but it is not clear how sensitive this number is to key parameters like the discount factor β, the utility premium from homeownership θ, and the parameters that drive the rental supply elasticity. Given that the welfare conclusions are the paper's most policy-relevant output, a sensitivity analysis over these parameters would substantially strengthen the reader's confidence in the 1.5% figure.

   (e) It would also be informative to decompose the welfare cost by age cohort at the time of the reform and not just for newly born households since the transitional costs fall disproportionately on cohorts who are already renters when the reform hits and face both higher rents and tighter borrowing limits simultaneously.

(f) Figure 7 Panel B shows that the welfare effect for homeowners turns slightly positive after age 40, but the paper does not explicitly rationalize this result in the text. The most natural explanation is that 40-year-old homeowners have accumulated enough wealth to enter landlordship at depressed house prices while having enough remaining lifetime to collect the present value of permanently higher rents. A brief explanation in the text would improve clarity.

(g) The maximum property ownership cap of three units is a computational simplification that the authors acknowledge. It is worth clarifying how sensitive the rental supply elasticity (and therefore the quantitative rent response) is to this cap. If allowing four or five units meaningfully changes the elasticity, the quantitative conclusions could be affected even if the qualitative mechanism is unchanged.