Mortgage Borrowing and the Boom-Bust Cycle in Consumption and Residential Investment

Xiaoqing Zhou*

Federal Reserve Bank of Dallas

November 2020

Abstract

This paper studies the transmission of the major shocks in the U.S. housing market in the 2000s to consumption and residential investment. Using geographically disaggregated data, I show that residential investment is more responsive to these shocks than consumption, as measured by elasticities and the implied contributions to GDP growth. I develop a structural life-cycle model featuring multiple types of housing investment to understand the large responses of residential investment. Consistent with the microdata, the model generates lumpy debt accumulation, lumpy housing investment and a strong correlation between mortgage borrowing and housing investment at the early stage of the life cycle. In the model, households move up the property ladder by increasing their mortgage debt after they have accumulated enough home equity. Since liquidity constraints and fixed costs prevent especially young homeowners from acquiring their desired home, shocks to their borrowing capacity have a large impact on residential investment.

Keywords: Mortgage borrowing, Consumption, Residential investment, House price, Mortgage rate, Credit supply, Business cycle.

JEL Codes: D1, E2, E3.

Declarations of interest: None.

*Federal Reserve Bank of Dallas, 2200 N. Pearl St., Dallas TX 75021. Email: xqzhou3@gmail.com. I thank Jason Allen, David Berger, Michael Gelman, Adam Guren, Joshua Hausman, Matteo Iacoviello, Lutz Kilian, Frank Stafford, Dmitriy Stolyarov, Joseph Vavra, Eleanor Wilking and participants at various seminars and conferences for helpful comments and discussions. An earlier version of this paper was completed while the author was at the Bank of Canada. The views in this paper are solely the responsibility of the author and should not be interpreted as reflecting the views of the Federal Reserve Bank of Dallas, the Federal Reserve System, or the Bank of Canada.

1 Introduction

In the years leading up to the 2008 financial crisis, the U.S. housing market experienced sharply rising house prices, declining mortgage rates and an expansion of credit. As a result, households aggressively borrowed against their homes and household debt reached an unprecedented level. Mortgage borrowing by existing homeowners has been widely viewed as instrumental for the propagation of housing market shocks to the broader economy in the 2000s. Since household borrowing is closely related to spending, an important question is to what extent shocks to the housing market affect aggregated household expenditures. To date, the literature has focused on the response of consumption to these shocks. The response of residential investment, in contrast, has received little attention. As Figure 1 shows, residential investment, like consumption, co-moved strongly with mortgage borrowing during the last housing cycle.

In this paper, I explore empirically and theoretically how consumption and residential investment responded to the major shocks in the U.S. housing market in the 2000s. While my analysis corroborates existing studies on the consumption responses, it shows that residential investment was more responsive to these shocks than consumption, as measured by elasticities and the implied contributions to GDP growth. This result holds both in the data and in a structural model that is consistent with the micro-level evidence on collateralized borrowing and household expenditures. Thus, residential investment played a key role in the propagation of housing market shocks to aggregate expenditures during the last housing cycle.

The first part of the paper provides detailed empirical evidence on the responses of consumption and residential investment to three major housing market shocks: shocks to house prices, mortgage rate shocks and credit conditions. Identifying the causal effects of these shocks is known to be difficult. I employ empirical strategies recently developed in the literature to identify each shock using geographically disaggregated data. To facilitate the comparison across shocks and outcomes, I use the same measures of expenditures at the metropolitan statistical area (MSA) level for all shocks.

There are three key empirical findings. First, the magnitude of the consumption responses is consistent with estimates in the existing literature. Second, the response of residential investment to each of these shocks is much larger than that of consumption. Third, even after taking into account the much lower share of residential investment in aggregate expenditures, the implied contribution of residential investment to output growth is larger than that of consumption. For example, in response to a one standard-deviation shock to the house price, all else equal, consumption contributes a 0.3 percentage point (pp) increase to annual GDP growth, while residential investment contributes a 0.4 pp increase. In response to mortgage rate or credit supply shocks, the contribution of residential investment relative to that of consumption is even larger. These results jointly suggest that residential investment was as important as consumption, if not more important, in driving fluctuations in U.S. aggregate expenditures in the 2000s.

This raises the question of whether a standard macro model with consumption and housing investment can explain this empirical evidence. If it does, what mechanisms give rise to the large responses of residential investment at the aggregate level? More importantly, is the model supported by the micro-level evidence on the link between borrowing and expenditures, as well as the heterogeneity of this link across households? To answer these questions, I develop a quantitative life-cycle model building on the workhorse consumption model of Berger et al. (2018). This type of model has been shown to be able to capture the heterogeneous consumption responses as in the microdata. The novel feature of my model is that households can choose between different types of housing investment (moving to a new home or making home improvements). This model feature allows me to quantify the use of mortgage borrowing for different kinds of housing investment as in the data, and to study which type of investment is more responsive to shocks.

The other elements of the model are more standard. Over the life cycle, households earn income, accumulate assets, and make spending and borrowing decisions. There are two consumption goods: non-housing and housing, and three assets: liquid savings, housing, and mortgage debt. Households face uninsurable labor income risks, collateral constraints and liquidity constraints. Mortgage debt contracts are long-term, and deviating from the contracted debt level is costly. These features ensure that the model generates rich heterogeneity in household balance sheets, consumption and housing investment, and that the model can be directly evaluated by microdata.

Data simulated from the steady state of the model match the salient life-cycle and cross-sectional dimensions of microdata. In particular, the model generates a strong correlation between homeowners' mortgage borrowing and their housing investment, both at the extensive and intensive margins, indicating the substantial use of mortgage loans by existing homeowners for housing investments (either through home equity extraction or new home-purchase loans). Moreover, these correlations are strongest among young homeowners. These patterns arise in the model, because young households start their life with an affordable home, the size of which is below the target level.

Thus, additional investments are desirable. However, the income of young households is too low and they do not have much liquid savings. As a result, households accumulate their home equity by gradually paying down the current mortgage and making small housing investments. When they have accumulated enough home equity, they pay a fixed cost to tap this equity and reinvest it in housing in a lumpy way. This explains the correlation between borrowing and housing investment in the data, and illustrates how households use debt to move up the property ladder in the absence of aggregate shocks.

When shocks to the housing market increase the borrowing capacity of households (e.g., higher house prices, lower mortgage rates and easier access to credit), households push forward their housing investments or increase the amount they would have invested otherwise, which leads to large responses of residential investment at the aggregate level. Using the model, I quantify the impacts of the three housing market shocks on consumption and residential investment, and show that they are consistent with the empirical evidence in Section 3. The model suggests that each shock spurs the borrowing of young households on average, and that the spending on housing investment increases most for those who raise their mortgage debt.

In the last part of the paper, I consider the joint impact of the three housing market shocks in the context of the U.S. economy in the 2000s. For this purpose, I simulate the evolution of aggregate consumption and residential investment by feeding the time series of real house prices, real mortgage rates, and a measure of credit conditions into the model. The simulated paths match broadly the boom-bust cycles in the corresponding U.S. data. Decomposing the two aggregate series suggests that housing investments by young homeowners disproportionately accounted for residential investment fluctuations over this period, whereas their consumption does not appear to have driven aggregate consumption dynamics. In other words, residential investment is more closely related to the changing borrowing capacity of young households in the 2000s than consumption.

To formalize this policy implication, I conduct a counterfactual analysis and assess the extent to which volatility in aggregate expenditures would have been reduced if households' borrowing ability had been restricted through higher down-payment requirements. The results show that the volatility of both consumption and residential investment growth would have declined, and that the decline would have been larger for residential investment growth. For example, if households had been required to provide at least 50% of their house value as a down payment in 2000 (compared to the actual down payment requirement of 20%), the volatility of consumption growth in the 2000s would have fallen by 28%, whereas the volatility of residential investment growth would have fallen by 58%.

Relation to the Literature

This paper is closely related to a burgeoning literature using microdata and heterogeneous-agent models to understand the impact of the major housing market shocks on the U.S. economy through consumer spending. This literature has emphasized heterogeneous consumption responses to house price shocks (see Mian and Sufi (2011, 2014), Mian et al. (2013), Cloyne et al. (2019), Aladangady (2017), Berger et al. (2018), Kaplan et al. (2019) and Guren et al. (2020)), to mortgage rate shocks (see Di Maggio et al. (2017), Bhutta and Keys (2016) and Chen et al. (2020)), and to monetary policy shocks (see Wong (2019), Greenwald (2018), Garriga et al. (2017), Garriga and Hedlund (2020), Cloyne et al. (2020), Beraja et al. (2019) and Berger et al. (2020)). My paper shows that these shocks exerted equally large or even larger effects on residential investment, which contributed to fluctuations in aggregate expenditures in the last housing cycle.

Another related literature empirically assesses the impact of credit supply shocks on bank lending, house prices and labor markets (see Loutskina and Strahan (2015) and Di Maggio and Kermani (2017)). My paper complements their findings by providing evidence on large residential investment responses. Consistent with the recent theoretical literature on the consumption response to these shocks (see Sommer et al. (2013) and Kaplan et al. (2019)), I find that, in the data and in my structural model, credit expansion has a limited effect on consumption.

On the household side, my paper is not the first to document the correlation between mortgage borrowing and housing investment (see Canner et al. (2002), Brady et al. (2000), Cooper (2009), Greenspan and Kennedy (2007), Benito and Power (2004) and Smith (2010)), but it is the first to document the heterogeneity of this correlation by age, and the first to use a structural life-cycle model to rationalize this correlation. Moreover, by modeling home upgrading and home improvements as alternative housing investment choices, my paper complements the work by Benmelech et al. (2019), who measure home improvements and related expenditures associated with home purchases, and explore the aggregate implications of this spending channel.

Finally, this paper relates to the DSGE model literature on the relation between collateralized borrowing and residential investment dynamics (see Khan and Rouillard (2018) and Iacoviello and Neri (2010)). My structural model complements theirs by introducing both the extensive and intensive margins of borrowing and investment, and by stressing the demographic aspect of the shock transmission, consistent with the microdata.

The remainder of the paper is organized as follows. Section 2 describes the data used for empirical analysis. Section 3 provides empirical evidence on the responses of consumption and residential investment to three major housing market shocks. Sections 4 and 5 outline the model and its calibration. Section 6 examines the cross-sectional and life-cycle patterns of borrowing and expenditures in the steady-state of the model and the corresponding patterns in the microdata. Section 7 analyzes the transmission of each shock through the lens of the model. Section 8 examines the boom-bust cycles in consumption and residential investment in the 2000s simulated from the model and conducts a counterfactual analysis. Section 9 concludes.

2 Data

The empirical analysis in this paper draws on a number of aggregate, regional, mortgage loan-level, and household-level data. It is useful to review the sources of these data and the key information contained in these datasets.

MSA-level consumption, residential investment and house price data. To measure consumption at the MSA-level, I follow Guren et al. (2020) in using retail employment in per capita terms between 1999 and 2015 from the Quarterly Census of Employment and Wages (QCEW). Given the scarcity of high-frequency, geographically disaggregated consumption data in the U.S. historically, Guren et al. provide evidence that retail employment has been viewed as one of the best proxies for consumer expenditures, and that it captures the variation in consumer expenditures nearly one-for-one both in the aggregate and in cities for which a direct consumption measure is available. MSA-level residential investment is measured by the valuation of building permits required for new single-family housing units, obtained from the Census Bureau Building Permits Survey. At the aggregate level, building permits capture most of the variation in residential investment. The correlation between the two time series is 0.92 for the period of 1999-2015. To measure changes in house prices at the MSA level, I use Freddie Mac House Price Indices (FMHPI) as the primary source and conduct robustness checks using alternative house price indices published by the Federal Housing Finance Agency (FHFA).

Loan-level mortgage origination and servicing data. For the analysis of mortgage rate shocks, I use the Residential Mortgage Servicing Database from Black Knight Financial Services (Black Knight McDash Data). The key to identification is measuring the incentive to prepay driven by the gap between a borrower's current mortgage rate and the prevailing market rate. The advantage of this dataset is its comprehensive and representative coverage of outstanding mortgage loans in the United States.[1] In my analysis, these loan-level data are used to construct three quarterly MSA-level variables: (i) the fraction of loans that have a positive rate gap at the beginning of each quarter, (ii) the average rate gap across all loans at the beginning of each quarter, and (iii) the fraction of loans that are prepayed voluntarily in the following year. In constructing these variables, I restrict the sample to first-lien 30-year fixed-rate mortgages in the McDash data for 1999-2015.

Loan-level mortgage application and approval data. For the analysis of credit supply shocks, I use the Home Mortgage Disclosure Act (HMDA) data for 1994-2006. This loan-level dataset covers the vast majority of home mortgage applications and approvals in the United States. It contains detailed information on loan and borrower characteristics at the time of the loan application. Two key variables are constructed using the HMDA data. One is the total number of mortgages originated in a given year at the MSA level. The other is the fraction of mortgage applications in a given MSA that are within 5% of the conforming loan limit in a year, which measures the exposure of an MSA market to an unexpected change in the conforming loan limit.

Household survey data. To evaluate whether the proposed structural model is able to match the cross-sectional heterogeneity in household expenditures associated with rising mortgage debt, I use data from the Panel Study of Income Dynamics (PSID) for 1999-2015. The sample is restricted to existing homeowners between age 26 and 65 who do not own farms or businesses. The data are used to construct four sets of variables for the empirical analysis in Section 6: (i) an indicator for a homeowner increasing her total mortgage debt, (ii) and indicator for a homeowner making housing investment, (iii) the real expenditures on housing investment and consumption, and (iv) household-level characteristics. A detailed description of the data and the variable definitions can be found in Appendix B.

3 Empirical Evidence

In this section, I estimate the responses of consumption and residential investment to three major housing market shocks: (i) a house price shock, (ii) a mortgage rate shock, and (iii) a credit supply shock. I employ empirical strategies recently developed in the literature to identify each type of

[1] The Black Knight McDash data cover approximately two-thirds of installment-type loans in the residential mortgage servicing market and contain more than 150 million individual loans with monthly performance history since April 1992.

shock using geographically disaggregated data. To facilitate the comparison across shocks and outcomes, my analysis uses the same measures of expenditures at the MSA level for all shocks.

3.1 House Price Shocks

Fluctuations in house prices have been one of the primary drivers of household borrowing and expenditures. Whereas a large body of empirical work has focused on the impact of house price shocks on consumption, little evidence has been provided for the response of residential investment at disaggregated levels.[2] With MSA-level data, these spending responses can be estimated by the following specification,

$$\Delta y_{m,t} = \gamma_t + \zeta_m + \beta_1 \Delta hp_{m,t} + \mathbf{x}_{m,t}\beta_2 + \varepsilon_{m,t}, \tag{1}$$

where $\Delta y_{m,t}$ denotes the log annual change in consumption or residential investment expenditures in MSA m at time t. $\Delta hp_{m,t}$ denotes the log annual change in house prices. γ_t is the time fixed effect controlling for the aggregate time trend, and ζ_m is the MSA fixed effect that captures unobserved heterogeneity. $\mathbf{x}_{m,t}$ denotes a set of additional control variables. The key parameter of interest, β_1, measures the response of consumption or housing investment to a 1% increase in the house price.

The OLS estimates are likely to be biased, as extensively discussed in the literature, mainly due to potential omitted variables that cause the error term, $\varepsilon_{m,t}$, to be correlated with house price growth (e.g., changes in income itself or the expectations about future income growth). To achieve identification, I employ two alternative instrumental-variable (IV) approaches recently proposed in the literature. The first approach is interacting the Saiz housing supply elasticity with national house price growth to instrument for MSA-level house price changes. Exogenous variation comes from the fact that when a housing boom appears nationwide, cities with inelastic housing supply will experience higher house price growth than cities with elastic housing supply.[3] This approach is implemented as a two-stage least square (2SLS) estimator. The first stage is in equation (2), and the second stage is specified as in equation (1):

$$\Delta hp_{m,t} = \gamma_t + \zeta_m + \alpha_1 \Delta nhp_t \times Saiz_m + \mathbf{x}_{m,t}\alpha_2 + u_{m,t}, \tag{2}$$

[2] A few studies have used microdata from the 1990s to estimate the effect of LTV changes, mostly driven by changes in house prices, on household mobility (see Genesove and Mayer (1997), Henley (1998) and Chan (2001)).

[3] Compared to using the Saiz elasticity alone as the instrument in a single cross-section, this interaction approach has the advantage that city-specific long-run trends correlated with the elasticity measure can be absorbed by the fixed effects in the panel-data setting (see Guren et al. (2020) and Davidoff (2016)).

where Δnhp_t denotes the log annual change in the national house price and $Saiz_m$ denotes the housing supply elasticity measure developed by Saiz (2010).

The second approach involves interacting the Guren et al. (2020) local house price sensitivity measure with national house price growth to instrument for MSA-level house price changes.[4] Guren et al. show that this sensitivity measure captures a more comprehensive set of determinants of housing supply than the Saiz elasticity measure and can be used to construct a more powerful instrument for house price changes. The regression further includes census-division-by-time fixed effects and a rich set of controls for the possibly heterogeneous effects of other aggregate shocks. Specifically, the following 2SLS specification is estimated,

$$\begin{aligned} \textit{1st stage:} \quad \Delta hp_{m,d,t} &= \xi_d \times \gamma_t + \zeta_m + \alpha_1 \Delta nhp_t \times GMNS_m + \mathbf{x}_{m,d,t}\alpha_2 + u_{m,d,t} \\ \textit{2nd stage:} \quad \Delta y_{m,d,t} &= \xi_d \times \gamma_t + \zeta_m + \beta_1 \widehat{\Delta hp_{m,t}} + \mathbf{x}_{m,d,t}\beta_2 + \varepsilon_{m,d,t}, \end{aligned} \tag{3}$$

where subscript d denotes the census division where MSA m is located. $GMNS_m$ denotes Guren et al.'s local house price sensitivity measure. $\zeta_d \times \gamma_t$ is the census-division-by-time fixed effect. Following Guren et al. (2020), $\mathbf{x}_{m,d,t}$ includes three sets of control variables: (i) the differential sensitivity of MSA-level retail employment to regional retail employment, (ii) the differential sensitivity of MSA-level retail employment to real 30-year mortgage rates and to the Gilchrist and Zakrajsek (2012) measure of excess bond premia, and (iii) two-digit SIC industry shares.[5] The same sets of control variables are also included in the Saiz-instrument approach and the OLS estimation. All specifications are estimated using quarterly data from 1999 to 2015. Standard errors are clustered at the MSA level.

Table 1 shows the OLS and IV estimates. The estimates are broadly consistent across specifications. The OLS estimates show that a 1% increase in house prices raises consumption by 0.09%, in contrast to a 1.23% increase in residential investment. In line with previous studies, the consumption response based on the Saiz instrument tends to be larger than the OLS estimate, but smaller based on the Guren et al.'s sensitivity instrument. Since the latter instrument is more powerful (first-stage R^2 of 0.82), these estimates are preferable. Thus, a 1% increase in house prices

[4] I use national house price growth to construct the instrument (as opposed to using census-division house price growth), because division-level house price indices are not available in the FMHPI data. Instead of averaging individual MSAs' house price indices (as implemented by Guren et al.), my approach exploits variation in local house price growth driven by a national housing boom or bust.

[5] The differential sensitivities of MSA-level retail employment in (i) and (ii) are estimated by regressing $\Delta re_{m,d,t} = \xi_d + \theta_m \Delta Y_{d,t} + \epsilon_{m,d,t}$, and including $\hat{\theta}_m \Delta Y_{d,t}$ as a control variable in the 2SLS estimation, where $\Delta re_{m,d,t}$ denotes the log annual change in retail employment and $\Delta Y_{d,t}$ denotes the log annual change in an aggregate variable.

boost consumption by 0.07% and residential investment by 1.03%. Moreover, all three strategies reveal that residential investment is more responsive than consumption, with a relative magnitude of 14.7 in the preferred specification. Appendix A presents additional results based on alternative residential investment measures and house price indices, as well as sub-sample estimates, all of which support the robustness of the estimates in Table 1.

These elasticities can be converted to changes in spending out of housing wealth (referred to as MPCs in Table 1), which implicitly takes into account the relative share of these expenditures. For this purpose, I divide the estimated consumption elasticity by the average ratio of housing wealth to consumption and obtain an MPC of 3.6 cents, in line with existing estimates.[6] Likewise, given the average ratio of housing wealth to residential investment, the spending on housing investment out of a $1 of housing wealth is 4.5 cents.[7] Table 1 summarizes the MPCs under each approach and shows that the change in housing investment expenditures is as large as the change in consumption, even after accounting for the relatively small share of residential investment in overall expenditures.

The large and positive response of residential investment to a house price shock raises the question of what the channels are through which house prices affect investment on housing and which households are most responsive in making these investments. Answering these questions requires more detailed information about consumers, which may be obtained from household survey data. The drawback of using survey data, however, is less precise measurement and the difficulty of finding powerful instruments. Having said that, Appendix Table B2 provides suggestive evidence from PSID data. I estimate the heterogeneous effect of house price changes on household expenditures and borrowing. The sample is restricted to existing homeowners. There are three key findings. First, a positive house price shock has a large and significant effect on the likelihood of making a housing investment (moving to a bigger home or making home improvements). Second, the housing investment response is strongest among young homeowners and declines with the age of the homeowner. Third, young homeowners are most likely to increase their mortgage debt in response to higher house prices, and those who increased their debt level are most likely to make a housing investment. These facts jointly suggest that higher house prices stimulate housing investment

[6]For example, using aggregate data, Case et al. (2005, 2013) find an MPC out of housing wealth around 3 to 4 cents; Carroll et al. (2011) estimate a next-quarter MPC of 2 cents and the eventual effect of 9 cents. Using cross-sectional data, Mian et al. (2013) and Mian and Sufi (2014) estimate the MPC between 5 and 7 cents; Aladangady (2017) estimate an MPC of 4.7 cents for homeowners; Guren et al. (2020) estimate an MPC of 2.4 cents, and of 3.3 cents if the pre-1990 period is dropped.

[7]Housing wealth is measured as the market value of owner-occupied real estate from the Flow of Funds. The average ratio of housing wealth to consumption (personal consumption expenditures less housing services) between 1980 and 1998 is 1.93. The average ratio of housing wealth to residential investment between 1980 and 1998 is 23.1.

through the borrowing capacity of young homeowners who tend to be liquidity constrained and eager to move up the property ladder.

3.2 Mortgage Rate Shocks

Mortgage rate changes are considered another key driver of household spending. Identifying this effect is empirically challenging, because consumers in the U.S. can choose whether and when to prepay their mortgages (for refinancing or purchasing a new home), and the ability to obtain lower rates depends on the borrower's financial condition and creditworthiness (Beraja et al. (2019)).

The recent study by Berger et al. (2020) proposed a method to identifying cross-sectional variation in the mortgage rate.[8] Their approach exploits two facts. First, the vast majority of mortgages in the U.S. are long-term fixed-rate, and for many consumers, their current mortgage rate differs from the rate they would obtain for a new loan covering the same balance, i.e., the prevailing market rate. Second, consumers are more likely to prepay their mortgages as the gap between their current rate and the prevailing market rate widens. These two facts suggest that borrower-specific rate gaps are predetermined and can be used to predict prepayment and spending patterns.

Aggregating borrower-specific rate gaps to the MSA level, however, is non-trivial. Berger et al. (2020) show that the distribution of rate gaps across individuals matters for predicting aggregate prepayment rates, and that the fraction of loans with positive rate gaps is a better measure of this distribution than the average gap.[9] I therefore use this fraction in an MSA at the beginning of a quarter, denoted by $(frac > 0)_{m,t-1}$, as the preferred measure of mortgage rate shocks, and estimate

$$\Delta y_{m,t} = \gamma_t + \zeta_m + \beta_1 (frac > 0)_{m,t-1} + \mathbf{x}_{m,t}\beta_2 + \varepsilon_{m,t}. \tag{4}$$

The regression includes MSA and quarter fixed effects, as well as a set of control variables. The key parameter of interest, β_1, measures the effect of having one additional percent of loans with positive rate gaps.

Panel I of Table 2 shows the effects on consumption and residential investment relative to four

[8]Alternative empirical strategies include focusing on adjustable-rate mortgage borrowers in a quasi-experimental design (Di Maggio et al. (2017)), using a large number of control variables to isolate exogenous changes in the benchmark mortgage rate (Bhutta and Keys (2016)), and estimating the impact of monetary policy shocks by household mortgage status (Wong (2019) and Cloyne et al. (2020)).

[9]Specifically, Berger et al. (2020) show that the likelihood of prepayment exhibits a step-like pattern: It is low and stable for negative rate gaps, then increases sharply as the gap becomes positive, and stabilizes again when the gap exceeds certain level.

quarters earlier. A 1 pp increase in the fraction of loans with positive gaps raises consumption by about 0.006% in a year. This effect is statistically indistinguishable from zero. In contrast, the same shock leads to a 0.5% increase in residential investment that is significant at the 1% level. The response of residential investment is likely driven by existing homeowners taking advantage of lower rates to finance their housing investment, for example, by originating a new home-purchase loan or making the down payment for another property through a cash-out refinance. The last column of panel I provides some evidence for this interpretation by estimating the change in the voluntary prepayment rate in the following year. It shows that a 1 pp increase in $(frac > 0)_{m,t-1}$ raises the prepayment rate in the following year by 0.2 pp, corresponding to a 15% increase in the annual prepayment rate, consistent with the estimates in Berger et al. (2020).

Appendix A provides additional results using alternative residential investment measures and an expanded set of controls, as well as sub-sample estimates, all of which support the large and positive response of residential investment. It is worth noting that, in the post-recession period, consumption responded positively and significantly to a mortgage rate shock, consistent with the evidence in Di Maggio et al. (2017) and Beraja et al. (2019) that the aggressive steps taken by the U.S. Federal Reserve during the crisis effectively brought down mortgage rates and stimulated household consumption through the refinancing channel.

Although $frac > 0$ better captures the distribution of borrower-level rate gaps in predicting aggregate prepayment rates, it is useful to provide estimates that can be interpreted as semi-elasticities, i.e., the responses to a 1 pp decline in the mortgage rate. It can be shown that the average rate gap at the MSA level, labeled as $Gap_{m,t}$ is highly correlated with $(frac > 0)_{m,t}$. In some cases, however, changes in $Gap_{m,t}$ are not reflected in $(frac > 0)_{m,t}$, which implies that the average rate gap is a noisy measure of the prepayment incentive at the aggregate level. One way of addressing this concern is to use $(frac > 0)_{m,t-1}$ as an instrument when estimating interest-rate semi-elasticities. Another complication is that the distribution of MSA-level average rate gaps has shifted to the right after the crisis: the mean increased from 45 basis points before 2008 to 119 basis points after 2008.[10] To enhance the predictive power of the average rate gap, I estimate a 2SLS specification that allows for a differential effect in the post-crisis period:

$$\Delta y_{m,t} = \gamma_t + \zeta_m + \beta_1 \widehat{Gap}_{m,t-1} + \beta_2 \widehat{Post_t \times Gap}_{m,t-1} + \varepsilon_{m,t}. \tag{5}$$

[10]Regressing $Gap_{m,t}$ on $(frac > 0)_{m,t}$ and $(frac > 0)_{m,t} \times \mathbb{I}(Post2008)$ shows a large differential effect for the post-crisis period (i.e., the coefficient of the interaction term), whereas regressing the prepayment rate on these two variables does not suggest a differential effect, validating the specification in equation (4).

where $Post_t$ is the indicator for the period after 2008. $\widehat{Gap}_{m,t-1}$ and $\widehat{Post_t \times Gap}_{m,t-1}$ are the predicted variables from the first-stage regression using $(frac > 0)_{m,t-1}$ and $Post_t \times (frac > 0)_{m,t-1}$ as the instruments.

Panel II of Table 2 shows the responses to a 1 basis point increase in the average rate gap. Consumption increases by 2.8 basis points, with an additional 1.2 basis points in the post-crisis period, consistent with the earlier evidence based on $frac > 0$. Residential investment increases by 62 basis points, and the effect is somewhat smaller after the crisis but still economically and statistically significant. The last column shows that these spending effects are likely realized through mortgage prepayments. To summarize, a fall in the mortgage rate boosts both residential investment and consumption (especially in the post-crisis period) through the mortgage prepayment channel. In particular, residential investment is much more responsive than consumption, as measured by semi-elasticities.

3.3 Credit Supply Shocks

Previous empirical studies have found that the expansion of mortgage credit in the early 2000s led to increased lending of financial institutions, raised house prices and stimulated the real economy (see Loutskina and Strahan (2015) and Di Maggio and Kermani (2017)). The purpose of this subsection is to quantify the effects of credit supply shocks on consumption and residential investment.

To identify these shocks, I follow Loutskina and Strahan (2015) in exploiting variation in the conforming loan limit specified by the GSEs (Fannie Mae and Freddie Mac) and in the heterogeneous exposure to these policy changes at the MSA level (see also Greenwald and Guren (2019)). This strategy utilizes the fact that the change in the conforming loan limit is exogenous to individual MSAs' economic conditions, and that borrowers applying for a loan amount close to this limit are funding-constrained. When the conforming loan limit is lifted, these borrowers are able to obtain cheaper credit, and MSAs with more constrained borrowers before the policy change tend to experience larger credit supply shocks.[11]

Specifically, two instrumental variables are constructed for the growth of mortgage credit. One is the faction of mortgage applications in an MSA that are within 5% of the conforming loan limit in the previous year, $AroundCutoff_{m,t-1}$. The second is the interaction between $AroundCutoff_{m,t-1}$ and the change in the conforming loan limit, $\Delta Cutoff_t$, that only varies over time. The 2SLS

[11]The GSEs can only purchase mortgages below the conforming loan limit. This limit increases each year by the percentage change in the national average of single-family median housing prices in the prior year.

regression is specified as

$$\begin{aligned} \textit{1st stage:} \quad \Delta Credit_{m,t} &= \alpha_1 AroundCutoff_{m,t-1} + \alpha_2 \Delta Cutoff_t \times AroundCutoff_{m,t-1} \\ &\quad + \mathbf{x}_{m,t}\alpha_3 + \gamma_t + \zeta_m + u_{m,t} \\ \textit{2nd stage:} \quad \Delta y_{m,t} &= \zeta_m + \gamma_t + \beta_1 \widehat{\Delta Credit}_{m,t} + \mathbf{x}_{m,t}\beta_2 + \varepsilon_{m,t}, \end{aligned} \tag{6}$$

where $\Delta Credit_{m,t}$ is the percent change in mortgage originations in MSA m and year t. The regression includes MSA and year fixed effects, as well as a set of control variables. The data are annual, due to the reporting period in the HMDA data, and extend from 1994 to 2006. The choice of the sample period follows Loutskina and Strahan (2015), who note that the instruments have become less powerful and potentially endogenous to local fundamentals since the onset of the crisis.

Table 3 presents the results. The IV estimates are larger than the OLS estimates. A 1% increase in mortgage credit raises residential investment by 0.9% (significant at the 1% level), but the effect on consumption is small, 0.023%, and insignificant. The latter evidence is new and interesting in its own right. The theoretical literature has long debated whether credit expansion can have a meaningful impact on consumption. Recent quantitative analysis suggests that this boils down to whether a large response of house prices is triggered by credit shocks (e.g., Kaplan et al. (2019)). Using the 2SLS approach, the house price response to a credit shock is estimated to be 0.14% (significant at the 5% level). Combining this estimate with the consumption elasticity with respect to the house price in Section 3.1, an alternative estimate of the consumption response to credit supply shocks is 0.01% (=0.14% ×0.07), close to the estimates in Table 3. These estimates do not appear to support a large impact of credit expansion on household consumption.

To ensure that the effects of credit supply shocks are not confounded by house price shocks and mortgage rate shocks, Appendix A presents the results including a set of additional controls. The consumption response is statistically insignificant, whereas the residential investment response is even stronger. Appendix A also provides other relevant robustness checks, all of which support the large and positive response of residential investment.

3.4 Summary

My empirical analysis shows that residential investment in geographically disaggregated data is much more responsive than consumption to each of the housing market shocks, when the responses are measured by elasticities (or semi-elasticities). The reason why most existing studies have focused on the consumption responses is perhaps that consumption accounts for a predominant share of GDP, whereas residential investment does not. I conclude this section by providing back-of-the-envelope calculations of the extent to which the consumption and residential investment responses have contributed GDP growth in the last housing cycle, taking into account their shares in aggregate expenditures. Table 4 shows the step-by-step calculation.

Although consumption accounts for a much larger share of GDP than residential investment (65% vs. 5% in 1999), the responsiveness of residential investment makes it at least as important a contributor to aggregate output fluctuations. For example, all else equal, a one standard deviation increase in the house price (7%) raises consumption by 0.49 pp, contributing a 0.32 pp increase to GDP. The same shock increases residential investment by 7.2 pp, translated to an increase in GDP by 0.36 pp. In response to mortgage rate or credit supply shocks, the contribution of the residential investment response to GDP growth relative to that of consumption is even larger. These results highlight the role of residential investment in driving aggregate fluctuations, which has been largely overlooked in the literature.

4 Model

Having empirically established that residential investment is more responsive than consumption to aggregate shocks and that it contributes more than consumption to output fluctuations, I develop a structural model to understand the mechanisms behind this finding. The model features rich heterogeneity on the household side, with a special focus on the decision of making housing investments over the life cycle. The micro-level predictions of the model are evaluated using household-level data in Section 6. The model is then used to understand the implications of micro-level decisions for aggregate expenditures in a boom and bust cycle (Sections 7 and 8).

The model builds on the quantitative life-cycle model of Berger et al. (2018). This type of model can capture the heterogeneity in the consumption responses in the microdata. The new feature of my model is that households can choose between different types of housing investment. Another key difference is that my model distinguishes between long-term debt and short-term liquid savings, allowing mortgage rates to affect household decisions separately from short-term interest rates.

Preferences. The economy is populated with overlapping generations of households who live for J periods. In the first J_y periods, households work. They are retired in the remaining $J - J_y$ periods. Households are endowed with an initial housing stock, a pre-existing mortgage balance, and an initial amount of liquid assets. They have preferences over consumption and housing services and have a bequest motive. A household born in t maximizes expected lifetime utility

$$\mathbb{E}_t \left[\sum_{j=0}^{J-1} \beta^j u(c_{j,t+j}, h_{j,t+j}) + \beta^J \Phi(w_{J,t+J}, h_{J,t+J}) \right],$$

where the first subscript of a variable denotes the household's age and the second subscript denotes time. The household index, i, is suppressed. c and h denote consumption and housing, respectively. $\Phi(\cdot)$ is the bequest function. $w_{J,t+J}$ is wealth at the end of life.

Income. Households face idiosyncratic income shocks during the working life. The process generating the logarithm of income is

$$\log(y_{j,t}) = \chi_j + z_{j,t}, \tag{7}$$

where χ_j is the deterministic component, invariant across households of the same age. $z_{j,t}$ is the idiosyncratic component, evolving according to

$$z_{j,t} = (1 - \rho_z)\bar{z} + \rho_z z_{j-1,t-1} + \varepsilon_{j,t}, \quad \varepsilon_{j,t} \sim^{i.i.d.} (0, \sigma_{ez}^2), \tag{8}$$

where $\bar{z}$ is the unconditional mean of z, ρ_z is the persistence parameter, and $\varepsilon_{j,t}$ is a mean zero i.i.d. shock with variance σ_{ez}^2. Households have rational expectations about their future income. Given $z_{j,t}$, $z_{j+1,t+1}$ is normally distributed with a conditional mean of $(1 - \rho_z)\bar{z} + \rho_z z_{j,t}$ and a conditional variance of σ_{ez}^2. At retirement, households face no income uncertainty and receive a fixed payment every period that is a fraction of the income in the final year of their working life.

Mortgage debt and liquid savings. Mortgages are long-term debt and are amortized over borrowers' remaining life.[12] The amortization schedule can be calculated by setting a fixed payment every period and a full repayment at the end of the term. For example, given the principal amount, b, the mortgage rate, r^b, and the term, T, the periodic payment, $M(b, r^b, T)$, is determined by

$$M(b, r^b, T) = \frac{r^b b}{1 - (1 + r^b)^{-T}}. \tag{9}$$

The mortgage balance changes when a borrower makes a contracted payment. In the example

[12] The existence of long-term debt to some extent dampens the contraction of aggregate spending when house prices fall substantially, because lenders cannot force underwater homeowners to deleverage under long-term contracts (see Ganong and Noel (2020), Berger et al. (2018), and Guren et al. (2020)).

above, the mortgage balance at the beginning of the next period becomes $(1+r^b)b - M$.

Borrowers can refinance their mortgage at any time by paying off the current balance and originating a new mortgage of size b' at a fixed cost, $f^b b'$. The collateral constraint must be satisfied at origination, so that the balance is less than a fraction of the home value:

$$0 \leq b' \leq \frac{(1-\theta)(1-\delta)}{1+r^b}\lambda ph.$$

This fraction is jointly determined by the minimum down-payment ratio (θ), the housing depreciation rate (δ), the mortgage rate (r^b), and the credit condition (λ). Households are also subject to liquidity constraints. They can save in liquid assets but cannot borrow without housing collateral. Let a_j be the amount of liquid savings at the beginning of age j and r^a be the corresponding interest rate such that $r^a < r^b$. The liquidity constraint can be written as $a_{j+1} \geq 0$.

Housing investments. Households can choose between two types of housing investment: home upgrading and home improvement.[13] Home upgrading allows a household to choose the size of the new home without a constraint, but the transaction cost is high. In contrast, home improvement incurs a smaller transaction cost but is constrained by the size that can be added to the current home. Home upgrading in the model reflects homeowners selling the current house and moving to a larger one, whereas home improvement mirrors renovating or remodeling the current home.[14]

This modeling choice implies that home upgrading is associated with larger expenditures but performed less frequently than home improvement. Three facts from the microdata support this feature. First, in the PSID data, the average spending on home upgrading is 1.6 times as much as home improvements. Conditional on households making these investments, spending on home upgrading is four times as much as home improvements (\$144k vs. \$32k). Second, home upgrading is associated with a larger increase in the number of main rooms than home improvements (0.96 vs. 0.13). Third, home upgrading occurs less frequently in a year than home improvements (4% vs. 11%). These facts suggest that the key differences between the two types of investment are the amount spent and the frequency of spending.[15] In Appendix D, I consider a canonical life-cycle

[13]In residential investment data, single- and multi-family structures and home improvements account for the largest shares, 46% and 30%, respectively, followed by commissions and ownership transfer costs (22%) and others (2%).

[14]I focus on investment in owner-occupied housing, because the recent literature emphasizes the dual role of housing as durable goods and illiquid assets. Investments in non-primary residences seen in the data may be given the interpretation of upgrading or improvement in the model, depending on the size of the investment. In the PSID data, this type of investment accounts for a small share of overall housing investment, and including them in either category does not change the cross-sectional or life-cycle investment patterns.

[15]I focus on the size of housing entering the utility function, as conventionally modeled in the literature. It is true that in practice, some home improvements are performed to suit particular tastes or lifestyle. One could consider

model that does not distinguish between home upgrading and home improvement, the predictions from which are compared with the model proposed in this section.

Recursive formulation. At the household level, the state variables include age (j), home size (h), mortgage balance (b), liquid savings (a) and income (y). The aggregate state, S, is characterized by the house price (p), the mortgage rate (r^b) and the credit condition (λ), i.e., $S \equiv (p, r^b, \lambda)$. In the model, when aggregate shocks hit the economy, households expect the realizations of the future aggregate state to remain constant at the current level.[16]

At age j, a household maximizes the expected lifetime utility by comparing the value of refinancing the mortgage, V_j^R, and not refinancing, V_j^N. The value function is given by $V_j = \max\{V_j^R, V_j^N\}$. If the household refinances the mortgage, it solves:

$$\begin{aligned} V_j^R(h,b,a,y;S) = \max_{h',b',a',c} \quad & u(c,h) + \beta \mathbb{E}_j \left[V_{j+1}(h',b',a',y';S') \right] \\ s.t. \quad c + a' + p\left[h' - h(1-\delta)\right] &= y + (1+r^a)a + b' - (1+r^b)b - f^b b' - \mathbf{F}ph' \\ 0 \le b' &\le \frac{(1-\theta)(1-\delta)}{1+r^b}\lambda p h \\ a' &\ge 0. \end{aligned} \tag{10}$$

The last term in the budget constraint, $\mathbf{F}ph'$, is the housing transaction cost depending on the type of investment. Specifically,

$$\mathbf{F} = \mathbb{I}\left\{h' \ge (1+\underline{h}^{up})h\right\} f^{up} + \mathbb{I}\left\{(1+\underline{h}^{imp})h \le h' < (1+\underline{h}^{up})h\right\} f^{imp} + \mathbb{I}\left\{h' < (1-\delta)h\right\} f^{down}. \tag{11}$$

In equation (11), $\mathbb{I}$ is the indicator function. The first term on the RHS is the transaction cost for home upgrading as a fraction of the home value, where $1+\underline{h}^{up}$ denotes the threshold for the new home size (relative to the home size in the previous period), beyond which $f^{up}ph'$ must be paid for the transaction. The second term is the transaction cost for home improvement as a fraction of the home value, where $1+\underline{h}^{imp}$ denotes the threshold of the home size, beyond which and below $1+\underline{h}^{up}$, $f^{imp}ph'$ must be paid. By construction, $f^{up} > f^{imp}$ and $\underline{h}^{up} > \underline{h}^{imp}$. The last term is the transaction cost for downsizing the home. When the household keeps the home size unchanged or

introducing idiosyncratic housing preference shocks to allow housing choices to depend on personal preferences. I leave this complication for future research.

[16] This is supported by the fact that house prices, mortgage rates and credit measures are highly persistent. To see this, the AR(1) coefficients of the real house price index, the real mortgage rate, and the ratio of outstanding mortgage over holdings of residential real estate is 0.99, 0.98 and 0.99, respectively, in quarterly aggregate data.

performs small maintenance work, i.e., $(1-\delta)h \le h^{'} < (1+\underline{h}^{imp})h$, no transaction cost is paid.

The household problem in (10) can also be reformulated as the maximum of the value of upgrading the home ($V_j^{U,R}$), the value of improving the home ($V_j^{I,R}$), the value of downsizing the home ($V_j^{D,R}$), and the value of performing small or no maintenance work ($V_j^{N,R}$), conditional on refinancing the mortgage. That is, $V_j^R = \max\{V_j^{U,R}, V_j^{I,R}, V_j^{D,R}, V_j^{N,R}\}$. The value of upgrading the home conditional on refinancing, for example, can be written as

$$
\begin{aligned}
V_j^{U,R}(h,b,a,y;S) = \max_{h^{'},b^{'},a^{'},c} \quad & u(c,h) + \beta \mathbb{E}_j \left[V_{j+1}(h^{'},b^{'},a^{'},y^{'};S^{'}) \right] \\
s.t. \quad c + a^{'} + p\left[h^{'} - h(1-\delta)\right] &= y + (1+r^a)a + b^{'} - (1+r^b)b - f^b b^{'} - f^{up} p h^{'} \\
h^{'} &\ge (1+\underline{h}^{up})h \\
0 \le b^{'} &\le \frac{(1-\theta)(1-\delta)}{1+r^b}\lambda p h \\
a^{'} &\ge 0.
\end{aligned}
$$

$V_j^{I,R}$, $V_j^{D,R}$ and $V_j^{N,R}$ are set up similarly with corresponding changes to the housing transaction cost and to the home size constraint.

If the household does not refinance the mortgage, it solves:

$$
\begin{aligned}
V_j^{N}(h,b,a,y;S) = \max_{h^{'},a^{'},c} \quad & u(c,h) + \beta \mathbb{E}_j \left[V_{j+1}(h^{'},b^{'},a^{'},y^{'};S^{'}) \right] \\
s.t. \quad c + a^{'} + p\left[h^{'} - h(1-\delta)\right] &= y + (1+r^a)a - M - \mathbf{F} p h^{'} \\
b^{'} &= (1+r^b)b - M \\
a^{'} &\ge 0,
\end{aligned}
$$

where $M \equiv M(b, r^b, J-j+1)$ is the scheduled mortgage payment as in formula (9). Similar to the case of refinancing, V_j^N can be reformulated as $V_j^N = \max\{V_j^{U,N}, V_j^{I,N}, V_j^{D,N}, V_j^{N,N}\}$, where $V_j^{U,N}$, for example, is the value of upgrading the home conditional on not refinancing, i.e.,

$$
\begin{aligned}
V_j^{U,N}(h,b,a,y;S) = \max_{h',a',c} \quad & u(c,h) + \beta \mathbb{E}_j \left[V_{j+1}(h',b',a',y';S') \right] \\
s.t. \quad c + a' + p\left[h' - h(1-\delta)\right] &= y + (1+r^a)a - M - f^{up} p h' \\
h' &\geq (1+\underline{h}^{up})h \\
b' &= (1+r^b)b - M \\
a' &\geq 0.
\end{aligned}
$$

Solution methods. The model is solved numerically using backward induction with a two-step procedure. In the first step, the state space is discretized and the value functions are solved over fixed grids of the state space. In the second step, policy functions are obtained by solving the optimization problem over finer grids, conditional on the value functions from the first step. The numerical procedure is described in detail in Appendix C.

5 Calibration

To assess the quantitative implications of the model, the model parameters are calibrated using external evidence and key moments in the microdata (see Table 5 for a summary). Households enter the life cycle at age 26, retire at the end of age 65, and die after 85, so $J = 60$ and $Jy = 40$. The initial home size, mortgage balance, and liquid assets are set to match their counterpart in the PSID data for households of age 21-25.

The utility function takes the Cobb-Douglas form, $u(c,h) = \left(c^{\alpha} h^{1-\alpha}\right)^{1-\sigma} / (1-\sigma)$, where σ is the inverse of the inter-temporal elasticity of substitution, and α is the expenditure share of consumption. Following the standard consumption literature, $\sigma = 2$. I set $\alpha = 0.81$ to match the share of housing investment expenditures in overall household spending in the PSID data, which is 12%. I set the discount factor $\beta = 0.935$ to match the age profile of the wealth-to-income ratio in the PSID data by minimizing the sum of squared distances between model-generated ratios and data-generated ratios in 5-year age bins. The bequest function is specified as $\Phi(w,h) = \eta(w^{\alpha} h^{1-\alpha})^{1-\sigma}/(1-\sigma)$, where $w \equiv (1+r^a)a - (1+r^b)b + ph(1-\delta)$. The bequest parameter, η, affects the borrowing decision only in the last few years of life. When η is small, the model generates large and increasing borrowing propensities toward the end of life, whereas the borrowing propensities in the data are small and non-increasing toward the end of life. I set $\eta = 4$, so that the borrowing propensity remains roughly constant during retirement.

I use the PSID data to estimate the income process. To obtain a smooth function for χ_j, I regress the log of deflated annual household income on the first- and second-order polynomials of age. I then calibrate ρ_z and σ_{ez} by taking the residual from the χ_j regression and estimating an AR(1) process, which gives $\rho_z = 0.9$ and $\sigma_{ez} = 0.18$.[17] At retirement, households receive a fixed payment every period that is a fraction, ζ, of their final-year permanent income. In the PSID data, annual household income of age 66 and above is 60% of the income of age 61-65, so ζ is set to 0.6.

The housing depreciation rate is set to $\delta = 0.0227$, following the estimate by the Bureau of Economic Analysis for residential 1- to 4-unit structures. The minimum down-payment ratio is set to $\theta = 0.2$, as home buyers seeking conventional financing would put down 20% of the home value. The mortgage refinancing cost, f^b, is set to 2%. The Federal Reserve Board's publication, *A Consumer's Guide to Mortgage Refinancings*, estimates the closing cost of refinanced mortgages to be 2-6% of the outstanding principal. I choose the lower end of this estimate because a higher value would result in a much lower rate of refinancing than in the microdata.

In the model, three transaction-cost parameters $(f^{up}, f^{imp}, f^{down})$ and two threshold-size parameters $(\underline{h}^{up}, \underline{h}^{imp})$ govern housing investment behavior. These parameters are calibrated as follows: (i) the transaction cost of home upgrading, f^{up}, is set to 2%, in line with the estimate by Freddie Mac that closing costs vary between 2% and 5% ($3,000 to $7,500) of the home value. A recent survey by Zillow shows that the average home buyer pays about $3,700 in closing fees, also implying a closing cost of 2%. (ii) Given f^{up}, the threshold size for home upgrading, $\underline{h}^{up}$, pins down the fraction of households who upgrade their homes. I set $\underline{h}^{up} = 15\%$ to match the annual rate of home upgrading in the PSID data, which is 4%. (iii) f^{down} governs the fraction of households who downsize their homes. I set $f^{down} = 3\%$ to match the annual downsizing rate in the PSID data, which is 1%. (iv) The remaining two parameters, f^{imp} and $\underline{h}^{imp}$, are calibrated jointly to match the fraction of households who make home improvements (11%), and the ratio of home upgrading expenditures to home improvement expenditures conditional on making these investments. This calibration results in $f^{imp} = 1\%$ and $\underline{h}^{imp} = 7\%$.

The steady state house price, p, is normalized to 1. The steady state mortgage rate r^b is set to 0.04, equal to the historical average of the 30-year fixed mortgage rate net of inflation. The return

[17] Since income is reported every other year in the PSID, equation (8) cannot be estimated directly. Iterating equation (8) once yields $z^i_{j,t} = (1-\rho_z^2)\bar{z} + \rho_z^2 z^i_{j-2,t-2} + (\varepsilon^i_{j,t} + \rho_z \varepsilon^i_{j-1,t-1})$, which allows for the estimation of ρ_z and σ_{ez}. Note that, in steady state, the average income at age j is equal to the permanent income, i.e., $\mathbb{E}y_j = \exp(\chi_j)$. Thus, $\mathbb{E}[\exp(z_j)] = 1$. Assuming z_j is normally distributed yields $\bar{z} = -\frac{\sigma_z^2}{2}$, where σ_z^2 is the variance of z_j.

on liquid assets, r^a, is set to 0.01, equal to the historical average of the 1-year treasury rate net of inflation. The measure for the credit condition in the steady state, λ, is set to 1.

6 Borrowing and Spending over the Life Cycle: Model and Microdata

In this section, I first use household survey data to establish a set of empirical facts about household borrowing and expenditures over the life cycle. I then show that the simulated data from the calibrated model can match these empirical facts quantitatively well. The model helps us to unpack the mechanisms behind these patterns, which will be important for understanding the response of residential investment to shocks at the aggregate level, as discussed in the next two sections.

Before I turn to data at the individual level, it is useful to examine life-cycle profiles for the average household. Figure 2 plots the paths of income, wealth (liquid and illiquid), consumption, housing investment and the likelihood of increasing one's mortgage debt (i.e., the borrowing propensity) simulated from the model's steady state and constructed using the PSID data. The model performs well in matching these profiles, including a hump-shaped path of income and consumption, rising wealth, and declining investment and borrowing propensities.[18] It shows that young households invest most in housing and are the most active borrowers.

6.1 Household Leverage, Housing Investment and Consumption

The life-cycle profiles for the average household, however, do not reveal whether those who increase their mortgage debt the most are also making the largest housing investments. Nor can we use them to quantify the change in expenditures associated with debt increases. To address these questions, I present three empirical facts from the PSID data that jointly characterize the relationship between household leverage, housing investment and consumption over the life cycle, and show that the simulated data from the model generate the same patterns. A detailed description of PSID data, key variables, and the additional analysis related to these facts can be found in Appendix B.

Fact 1. Young homeowners who increase their mortgage debt are most likely to invest in housing.

Homeowners may invest in their own housing by upgrading to a bigger home or improving the current home. Both actions, together with the change in mortgage debt, are observed in the data and in the model. These variables can be used to estimate the likelihood of making a housing

[18] As in related studies, the comparison focuses on the working-life period, since standard life-cycle models are not designed to capture income risks or health shocks during retirement that would be important for retirees' decisions.

investment associated with debt increases. To account for household characteristics, unobserved heterogeneity and aggregate shocks that are likely to confound the correlation of interest, I estimate the following logistic regression for the PSID data,

$$Pr(Invest_{i,t} = 1|Borrowed_{i,t}, \mathbf{x}_{i,t}, \gamma_t) = \mathbf{F}\left(\beta_0 + \beta_1 Borrowed_{i,t} + \mathbf{x}_{i,t}\beta_2 + \gamma_t\right), \qquad (12)$$

where i and t denote household and time. $Invest_{i,t}$ and $Borrowed_{i,t}$ are indicators for making housing investment and increasing mortgage debt. $\mathbf{x}_{i,t}$ is a set of controls for demographics (age, age-squared, and changes in family size), financial condition in the previous period (income, liquidity and illiquid assets), and changes in income, house prices and mortgage rates in the current period. γ_t is the time fixed effect.[19] For the simulated data, I estimate a similar specification that includes all controls except for changes in family size (not modeled) and changes in house prices and mortgage rates, as well as time fixed effects (not relevant for the steady state analysis).

Table 6 presents the estimates of β_1 using the PSID and the simulated data. Panel I shows that, in the microdata, when a household increases its mortgage debt, the likelihood of making a housing investment (either home upgrading or improvement) increases by 9 pp, corresponding to a 66% increase relative to the unconditional housing investment likelihood. Young homeowners, in particular, are most likely to use the borrowed funds for housing investment. Breaking down the housing investment into home upgrading and improvements, panels II and III show that increased household debt is correlated with a higher investment likelihood for each type. Along the age dimension, the correlation declines with age for home upgrading, whereas it exhibits a U shape for improvements (see detailed discussion in Section 6.2). The structural model performs quantitatively well in matching each of these patterns in the microdata.

Fact 2. Homeowners who increase their mortgage debt have a higher growth in housing investment expenditures, but not in consumption expenditures.

After showing that homeowners who increase their mortgage debt are more likely to invest in housing (i.e., the extensive margin effect), I now estimate the change in their housing investment expenditures relative to other homeowners (i.e., the intensive margin effect). For comparison, I also estimate the change in their consumption expenditures relative to other homeowners. In the PSID data, consumption is measured as the sum of non-durables, durables and services, while housing

[19] I do not include household fixed effects in logit regressions because consistent estimates may not be obtained (Wooldridge (2002)). However, I obtain very similar estimates when estimating alternative linear probability models with household fixed effects.

investment covers the expenditures on home upgrading and improvement, all converted to 2009 dollars. Using the microdata, I estimate the following specification

$$\frac{\Delta Exp^k_{i,t}}{TotalExp_{i,t-1}} = \zeta^k_i + \beta^k_1 Borrowed_{i,t} + \mathbf{x}_{i,t}\beta^k_2 + \gamma^k_t + \varepsilon^k_{i,t}, \tag{13}$$

where the LHS denotes the change in expenditure type k normalized by the household's total expenditures in $t-1$. The regression includes a set of control variables as in equation (12) and year and household fixed effects. For simulated data, a similar specification is estimated excluding the non-applicable control variables mentioned earlier.

Panel I of Table 7 shows the change in housing investment expenditures associated with debt increases. In the microdata, those who increase their mortgage debt on average see a 20% increase in housing investment expenditures relative to those who stay with their current mortgage contract. In the simulated data, this effect is 16%. Both the microdata and the model suggest that the change in housing investment expenditures financed by borrowing declines with age. In contrast, the change in consumption associated with higher debt is close to zero and statistically insignificant in the microdata (Appendix Table B3). Neither a break down by household age nor by consumption type suggests any noticeable heterogeneity.[20]

Panels II and III show the change in home upgrading and improvement expenditures associated with debt increases. The model captures the pattern that young households who increase their mortgage debt experience the largest increase in these expenditures. Moreover, both in the microdata and in the model, expenditures on home upgrading increase more than home improvements.

Fact 3. A $1 increase in mortgage debt is associated with 30-40 cents spending on housing investment.

An alternative measure of the intensive margin effect of household debt is the change in spending associated with a $1 increase in mortgage debt. Using the PSID data, this can be estimated by the following regression,

$$\Delta Exp^k_{i,t} = \zeta^k_i + \beta^k_1 \Delta b_{i,t} + \mathbf{x}_{i,t}\beta^k_2 + \gamma^k_t + \varepsilon^k_{i,t}. \tag{14}$$

[20] It is worth noting that the correlation between consumption growth and debt increases is weak for two reasons. First, absent aggregate shocks, consumption is roughly smooth over the life cycle. Second, the propensity to borrow is highest among young households, who primarily use the borrowed funds to invest in housing rather than for consumption. This steady-state pattern, however, does not contradict to the significant responses of consumption to aggregate shocks.

where $\Delta Exp_{i,t}^{k}$ and $\Delta b_{i,t}$ denote the dollar change in household i's expenditures and mortgage debt in period t. A similar specification is estimated for the simulated data excluding the non-applicable control variables mentioned earlier.

Table 8 shows that on average households spend 32 cents of borrowed funds on housing investment in the microdata, similar to the 39 cents in the model. A striking pattern is that, both in the microdata and in the model, young homeowners spend a predominant share of the borrowed funds on housing investment, and that this share declines sharply with household age. Consistent with the finding in Table 7, most of the funds are associated with expenditures on home upgrading not home improvements. The change in consumption, in contrast, is close to zero and does not exhibit significant heterogeneity across households (Appendix Table B3).

6.2 Understanding Household Investment Decisions

Why is the correlation between debt increases and housing investment so high, both at the extensive and intensive margins, especially for young homeowners? The model provides a structural explanation. Households start their lives with small houses, so the marginal utility of housing is high and additional investments are desirable. However, young households earn little income and do not have liquid savings to pay for these investments. As a result, they pay down the current debt and make small investments to build up their home equity. When they have accumulated enough home equity, they pay a fixed cost to tap this equity and reinvest it in housing, either by moving to a bigger home or by making home improvements. Because of transactions costs, investments are lumpy and are made at the time when households increase their mortgage debt, consistent with the patterns in the microdata.

An interesting question is how households choose between home upgrading and home improvement. In the absence of borrowing constraints and income uncertainty, young households would make large investment through home upgrading, which brings them closer to the target home size. In the presence of these frictions, the choice of investment depends on income and home equity. The model generates three predictions about the choice between home upgrading and home improvement. First, high income households, who are able to build up home equity quickly, choose home upgrading. Second, low income households, who cannot afford home upgrading and are slow to build up home equity, choose home improvement. Third, the correlation between home upgrading and debt increases is higher than the correlation between home improvement and debt increases, because home improvement spending is relatively low and is affordable for many households without incurring additional borrowing.

Figure 3 illustrates these predictions by plotting the life-cycle profiles of four individuals. The left column shows their realized income and timing of debt increases. The right column shows their home upgrading and improvement expenditures. Household (A) earns high income early in life, making it possible to upgrade the home in the first period and to reinvest the home equity in the subsequent period for another upgrading. Later in life, the household makes additional investment through home improvements without raising the debt level. Household (B) earns slightly lower income early in life, causing it to make several out-of-pocket investments before having enough home equity to upgrade its home. Household (C) earns relatively low income at the beginning of the life cycle, limiting the investment size and slowing down the accumulation of home equity. Only after the household has made several improvements, it is able to upgrade the home with the help of a larger mortgage. Household (D) earns low income early in life and cannot afford home upgrading. It ends up making a few home improvements, two of which are financed by debt increases.

The model predictions can be tested by the microdata. A straightforward testable implication is that the correlation between home upgrading and debt increases is stronger among high-income households, whereas the correlation between home improvements and debt increases is stronger among low-income households. Table 9 shows these correlations estimated using the PSID data and the simulated data. The microdata support this hypothesis and the empirical estimates are quantitatively similar to those simulated from the model.

7 Shocks and Aggregate Responses in the Model

This section examines the responses of consumption and residential investment to the three housing market shocks through the lens of the model. Simulations show that residential investment strongly comoves with borrowing and is more responsive to each shock than consumption, consistent with the empirical evidence. This conclusion, from the quantitative-theory point of view, is not new, as early DSGE models have explored this pattern (e.g. Iacoviello and Neri (2010)). The difference is that my model can answer the question of which households drive these responses and which type of investment is more responsive.[21]

There is one caveat. The analysis in this section examines one shock at a time, holding other aggregate variables fixed. While this helps to account for the direct impact, the partial-equilibrium

[21] In Appendix D, I present impulse responses simulated from my model and from Iacoviello and Neri (2010) model, and discuss in detail the similarities and the differences.

nature of this exercise means that it may miss important general-equilibrium effects. For example, lower mortgage rates and relaxed credit conditions are likely to change house prices, further affecting spending through this indirect channel. In Appendix F, I use the model to revisit the effect of mortgage rate and credit supply shocks with house prices moving simultaneously, and show that the responses are amplified to some extent but overall similar to those discussed in this section.

7.1 House Price Shocks

The upper panel of Figure 4 shows the impact responses of borrowing, housing investment and consumption to a 1% permanent house price increase. The borrowing propensity increases by 0.3 pp on average and declines with age, consistent with household-level evidence (Appendix Table B2). Young homeowners increase their housing investment because of relaxed collateral constraints and higher home equity, whereas old households cut back on this type of spending due to higher costs. Households also increase their consumption, particularly the young. As in Berger et al. (2018), a positive house price shock affects consumption through the collateral, wealth, income and substitution channels, all of which exert an effect of the same sign on housing investment, except for the substitution channel that dampens housing investment and stimulates consumption.

In the lower panel, the spending responses are shown for those who increase their debt ("increasers") and those who do not ("non-increasers"), as well as by investment type. This shows that the increase in home upgrading and improvement expenditures is mainly driven by young debt increasers who otherwise would not invest or who would invest less. Higher house prices boost their home equity immediately, allowing them to push forward or to increase their housing investment. Old homeowners, in contrast, see a decline in their investment because of the higher price of housing. These households, who would have invested mostly in home improvements, curtail these expenditures in response to the shock. In the meantime, they substitute consumption for housing, which explains their higher consumption response (dotted line, lower-right panel).

Figure 4 shows that higher house prices stimulate housing investment of young homeowners (mainly through the collateral and wealth channels), even though households do no expect future house prices to keep growing. Higher house prices may also change beliefs about future house prices, affecting spending through this additional channel (see Kaplan et al. (2019) and Bailey et al. (2019)). In Appendix G, I model the change in expectations driven by current house price shocks and show that the patterns in Figure 4 are largely reserved and even reinforced with this additional channel in operation.

7.2 Mortgage Rate Shocks

Mortgage markets are widely recognized as important for the monetary policy transmission. This is because, on the one hand, mortgage rates matter for consumer spending, and, on the other hand, they can be affected by monetary policy actions.[22] Garriga et al. (2017) provide a detailed quantitative analysis of the channels through which mortgage contracts transmit interest rate shocks. The two channels proposed in their analysis are at work in my model. First, a lower mortgage rate reduces the cost of new mortgage loans (i.e., refinanced mortgages in my model), generating a price effect for refinancing borrowers and stimulating their spending. Second, a lower mortgage rate brings down the payments on existing mortgages, generating an income effect for non-refinancing borrowers who hold adjustable-rate mortgages. The life-cycle dimension of my model implies that these effects should be stronger among young households.[23]

In the upper panel of Figure 5, the responses of borrowing, consumption, and housing investment to a 100 basis point reduction in the mortgage rate are plotted. Consistent with the empirical evidence, in the model, demand for mortgage refinancing increases, and young homeowners are more responsive than old homeowners. The borrowing response also suggests that the price effect is larger for homeowners around age 35 than for younger homeowners, because those around 35 have accumulated more home equity and can obtain larger mortgages in response to the shock. The age profiles of consumption and housing investment resemble that of the borrowing propensity.

The lower panel shows the spending responses of debt increasers and non-increasers. As expected, those who increase their mortgage debt benefit from lower costs of borrowing, and hence increase their housing investment and consumption. Interestingly, those who do not change their mortgage contract also increase their consumption and spending on home improvements. Since mortgages in my model are akin to the adjustable-rate mortgages in Garriga et al. (2017), this finding supports the income effect of rate declines. Note that the spending responses of non-increasers can be substantial relative to the current-period reduction in their mortgage payment. This is because some households (mostly middle-age), who otherwise would completely pay off their mortgages, choose to delay their prepayment and instead spend on housing and

[22]Recent quantitative studies have emphasized the role of the mortgage market in transmitting monetary policy shocks to the real economy. See, e.g., Garriga et al. (2017), Wong (2019), Greenwald (2018), Sommer et al. (2013), Garriga and Hedlund (2020) and Berger et al. (2020).

[23]Wong (2019) develops a quantitative life-cycle model to study the effect of monetary policy shocks on household consumption through the mortgage refinancing channel. The responses in her study, however, cannot be directly compared to the responses to a mortgage rate shock in the current paper, because monetary policy shocks in her model affect not only mortgage rates, but also other macroeconomic variables such as house prices and income.

consumption.

7.3 Credit Supply Shocks

Figure 6 shows the impact of a 1% permanent increase in the housing collateral value (λ). This shock generates similar effects to a positive house price shock, especially along the life-cycle dimension. There are two differences, however. First, the response of housing investment, especially of home upgrading, to the credit shock is larger than to a house price shock. This is because the credit shock, unlike the house price shock, does not change the relative price of housing to consumption and hence stimulates housing investment to a greater extent. Second, the consumption response to the credit shock is smaller than to the house price shock. This is due to the absence of the wealth and substitution effects that would further boost consumption to a house price shock.[24]

8 A Boom-Bust Episode

In this section, I use the model to examine whether the combination of housing market shocks that occurred in the 2000s can account for the fluctuations in aggregate consumption and residential investment, and which households played a key role in driving these fluctuations.[25] As stressed in the macro-housing literature (see Iacoviello and Neri (2010) and Davis and Heathcote (2005)), understanding residential investment dynamics is important for understanding business cycle dynamics, since residential investment is the most volatile of the major components of GDP, and as shown in the empirical evidence, its responses contributed disproportionately to GDP growth in the last housing cycle.

I simulate the evolution of aggregate consumption and residential investment by feeding the U.S. historical time series for real house prices, real mortgage rates, and credit conditions into the model (as shown in Figure H1). In particular, changes in the credit condition before the financial crisis are measured by the change in the GSE conforming loan limit. As secured lending fell since 2008, the post-crisis condition of credit is set to the level of 2000. Figure 7 shows the evolution of aggregate consumption, residential investment and its two components simulated from the calibrated model. Comparing these series with the actual data (Figure 1) suggests that the

[24]Recent quantitative studies with endogenous house prices have reached the same conclusion that credit supply shocks have a limited effect on consumption, precisely because these shocks do not change house prices materially (see Sommer et al. (2013) and Kaplan et al. (2019)).

[25]I focus on the period of 2000-2010 because after the crisis the government took various actions to regulate credit markets (e.g., Dodd-Frank), to stimulate home purchases (e.g., first-time homebuyer credit), and to relieve debt burdens of homeowners (e.g., HAMP and HARP). These policies introduced new shocks not captured by the model.

model matches key features of the boom-bust cycle in U.S. aggregate expenditures in the 2000s, both in the timing and the magnitude of the fluctuations.[26]

The model highlights an amplification channel for residential investment in the boom-bust cycle of the 2000s. In the first half of the decade, home prices appreciated strongly, credit conditions were relaxed and the mortgage rate fell. These developments stimulated existing homeowners to borrow aggressively. Young borrowers then spent most of the borrowed funds on housing investment, in particular, to upgrade their homes. As a result, the aggregate housing stock was built up, facilitating more borrowing and spending. In subsequent years, as these developments began to reverse, household wealth plunged, fewer homeowners were able to borrow, and spending on consumption and housing investment plummeted. Since housing investment at the individual level is more responsive to shocks than consumption, residential investment at the aggregate level is more volatile than aggregate consumption. As in the data, fluctuations in residential investment is mostly driven by home upgrading rather than home improvements, because young homeowners have higher demand for home upgrading and are also more sensitive to shocks.

To illustrate the importance of mortgage borrowing in driving the residential investment dynamics, Figure 8 plots the shares of consumption and housing investment accounted for by those who increased their mortgage debt. The consumption share of these borrowers is similar to their population share, reaching 26% in 2005 and falling below 10% after 2007. In contrast, their housing investment accounted for more than 80% of aggregate residential investment during the boom years and dropped to 20% during the financial crisis. Closer examination of the expenditure shares accounted for by young borrowers (age 25-45) who increase their mortgage debt shows that they were the driving force behind the fluctuations in residential investment, despite the fact that they only accounted for a small fraction of the entire population.

Finally, I conduct a counterfactual policy analysis to understand the role of rising household debt in driving fluctuations in aggregate expenditures. I ask how much lower the volatility of consumption and residential investment growth would have been, if households' ability to use their homes as collateral had been limited at the onset of the housing boom in 2000. To this end, I simulate the evolution of consumption and residential investment in the 2000s with the same time series fed into the model, but raising the minimum down-payment ratio, θ. For each value of θ, I

[26] A noticeable divergence between the model and the data is in 2006. Both consumption and residential investment were still rising in the data, whereas the model predicts the opposite. In the model, spending declines in 2006 because the mortgage rate rose sharply. This negative effect was not reversed by rising house prices or relaxed credit conditions. This suggests that rising spending in the data is likely driven by other shocks such as aggregate income shocks.

compute the standard deviation of expenditure growth as the volatility measure.

Figure 9 shows the volatility measures normalized by the corresponding values in the baseline model (i.e., $\theta = 0.2$). The main findings from this exercise can be summarized in two points. First, the volatility of each type of spending decreases with θ, as homeowners are unable to borrow as much as before and spend less. Second, as θ increases, the volatility of residential investment growth decreases faster than that of consumption. For example, if households have been required to provide at least 50% of the house value as the down payment in 2000, the volatility of residential investment growth in the 2000s would have fallen by 58%, whereas the volatility of consumption growth would have fallen by 28%. These findings complement the earlier conclusion that housing investment is more sensitive to shocks affecting the household borrowing capacity than consumption. Hence, restricting this capacity reduces fluctuations in residential investment to a larger extent.

9 Conclusion

In the first half of the 2000s, the U.S. experienced strongly rising home prices, falling mortgage rates and a credit expansion, along with a booming economy. In subsequent years, the reversal of these trends eventually pushed the economy into a severe recession. A large literature has focused on how these housing market shocks were transmitted to consumption through the borrowing of homeowners. To date, there has been no systematic analysis on the role of residential investment in driving aggregate fluctuations over this period, even though residential investment co-moved with mortgage borrowing and disproportionately contributed to GDP growth in that cycle. This paper addresses this question from an empirical, a theoretical and a policy point of view.

On the empirical side, this paper provides evidence based on geographically disaggregated data and state-of-art identification strategies that residential investment is more responsive to each of the three major housing market shocks than consumption, as measured by elasticities and the implied contributions to GDP growth. This evidence suggests that residential investment played at least as important a role as consumption in moving aggregate expenditures during the last housing cycle.

On the theoretical side, this paper develops a quantitative life-cycle model to help understand the mechanisms behind the large responses of residential investment. The model is consistent with the salient life-cycle and cross-sectional dimensions of the microdata. In particular, it replicates the use of mortgage debt by existing homeowners for housing investment, especially at the early age of the life cycle, as households move up the property ladder. The model shows that shocks that increase households' borrowing capacity push forward their housing investment decisions and enable them to invest more than they would have invested otherwise. This translates to large responses of residential investment at the aggregate level, together with surging household debt.

On the policy side, my analysis suggests that the boom-bust cycle in residential investment in the 2000s was mainly driven by the debt-financed investments of young homeowners. Restricting homeowners' borrowing ability at the onset of the housing boom would have greatly reduced the volatility of residential investment growth over this period.

References

Aladangady, A., "Housing Wealth and Consumption: Evidence from Geographically-Linked Microdata," *American Economic Review* 107(11) (2017), 3415–3446.

Bailey, M., E. Davila, T. Kuchler and J. Stroebel, "House Price Beliefs and Mortgage Leverage Choice," *Review of Economic Studies* 86(6) (2019), 2403–2452.

Benito, A. and J. Power, "Housing Equity and Consumption: Insights from the Survey of English Housing," *Bank of England Quarterly Bulletin* (2004), 302–309.

Benmelech, E., B. T. Melzer and A. M. Guren, "Making the House a Home: The Stimulative Effect of Home Purchases on Consumption and Investment," NBER Working Paper No. 23570, 2019.

Beraja, M., A. Fuster, E. Hurst and J. Vavra, "Regional Heterogeneity and the Refinancing Channel of Monetary Policy," *Quarterly Journal of Economics* 134(1) (2019), 109–183.

Berger, D., V. Guerrieri, , G. Lorenzoni and J. Vavra, "House Prices and Consumer Spending," *Review of Economic Studies* 85(3) (2018), 1502–1542.

Berger, D. W., K. Milbradt, F. Tourre and J. Vavra, "Mortgage Prepayment and Path-Dependent Effects of Monetary Policy," NBER Working Paper No. 25157, 2020.

Bhutta, N., "The Ins and Outs of Mortgage Debt during the Housing Boom and Bust," *Journal of Monetary Economics* 76 (2015), 284–298.

Bhutta, N. and B. J. Keys, "Interest Rates and Equity Extraction during the Housing Boom," *American Economic Review* 106(7) (2016), 1742–1774.

Brady, P., G. Canner and D. Maki, "The Effects of Recent Mortgage Refinancing," *Federal Reserve Bulletin* (2000), 441–450.

Canner, G., K. Dynan and W. Passmore, "Mortgage Refinancing in 2001 and Early 2002," *Federal Reserve Bulletin* (2002), 469–481.

Carroll, C. D., M. Otsuka and J. Slacalek, "How Large Are Housing and Financial Wealth Effects? A New Approach," *Journal of Money Credit and Banking* 43(1) (2011), 55–79.

Case, K. E., J. M. Quigley and R. J. Shiller, "Comparing Wealth Effects: The Stock Market vs. the Housing Market," *Advances in Macroeconomics* 5(1) (2005), 1–34.

———, "Wealth Effects Revisited: 1975-2012," *Critical Finance Review* 2(1) (2013), 101–128.

Chan, S., "Spatial Lock-in: Do Falling House Prices Constrain Residential Mobility?," *Journal of Urban Economics* 49(3) (2001), 567–586.

Chen, H., M. Michaux and N. Roussanov, "Houses as ATMs? Mortgage Refinancing and Macroeconomic Uncertainty," *Journal of Finance* 75(1) (2020), 323–375.

Christiano, L. J., M. Eichenbaum and C. L. Evans, "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," *Journal of Political Economy* 113(1) (2005), 1–45.

Cloyne, J., C. Ferreira and P. Surico, "Monetary Policy when Households have Debt: New Evidence on the Transmission Mechanism," *Review of Economic Studies* 87(1) (2020), 102–129.

Cloyne, J., K. Huber, E. Ilzetzki and H. Kleven, "The Effect of House Prices on Household Borrowing: A New Approach," *American Economic Review* 109(6) (2019), 2104–2136.

Cooper, D., "Did Easy Credit Lead to Overspending? Home Equity Borrowing and Household Behavior in the Early 2000s," FRB Boston Public Policy Discussion Paper, No. 2009-7, 2009.

Davidoff, "Supply Constraints Are Not Valid Instrumental Variables for Home Prices Because They Are Correlated With Many Demand Factors," *Critical Finance Review* 5(2) (2016), 177–206.

Davis, M. and J. Heathcote, "Housing and the Business Cycle," *International Economic Review,* 46(3) (2005), 751–784.

Di Maggio, M. and A. Kermani, "Credit-Induced Boom and Bust," *Review of Financial Studies* 30(11) (2017), 3711–3758.

Di Maggio, M., B. Keys, T. Piskorski, R. Ramcharan, A. Kermani, A. Seru and V. Yao, "Interest Rate Pass-Through: Mortgage Rates, Household Consumption, and Voluntary Deleveraging," *American Economic Review* 107 (2017), 3350–3388.

Ganong, P. and P. Noel, "Liquidity vs. Wealth in Household Debt Obligations: Evidence from Housing Policy in the Great Recession," *forthcoming: American Economic Review* (2020).

Garriga, C. and A. Hedlund, "Mortgage Debt, Consumption, and Illiquid Housing Markets in the Great Recession," *American Economic Review* 110(6) (2020), 1603–1634.

Garriga, C., F. Kydland and R. Sustek, "Mortgages and Monetary Policy," *Review of Financial Studies* 30(10) (2017), 3337–3375.

Genesove, D. and C. J. Mayer, "Equity and Time to Sale in the Real Estate Market," *American Economic Review* 87(3) (1997), 255–269.

Gilchrist, S. and E. Zakrajsek, "Credit Spreads and Business Cycle Fluctuations," *American Economic Review* 102(4) (2012), 1692–1720.

Greenspan, A. and J. Kennedy, "Sources and Uses of Equity Extracted from Homes," *Oxford Review of Economic Policy* 24 (2007), 120–144.

Greenwald, D. L., "The Mortgage Credit Channel of Macroeconomic Transmission," MIT Sloan Working Paper, 2018.

Greenwald, D. L. and A. Guren, "Do Credit Conditions Move House Prices," MIT Sloan School Working Paper 5910-19, 2019.

Guren, A. M., A. McKay, E. Nakamura and J. Steinsson, "Housing Wealth Effects: The Long View," *forthcoming: Review of Economic Studies* (2020).

Henley, A., "Residential Mobility, Housing Equity and the Labour Market," *Economic Journal* 108(447) (1998), 414–427.

Iacoviello, M. and S. Neri, "Housing Market Spillovers: Evidence from an Estimated DSGE Model," *American Economic Journal: Macroeconomics* 2 (2010), 125–164.

Kaplan, G., K. Mitman and G. L. Violante, "The Housing Boom and Bust: Model Meets Evidence," *forthcoming: Journal of Political Economy* (2019).

Khan, H. and J. Rouillard, "Household Borrowing Constraints and Residential Investment Dynamics," *Journal of Economic Dynamics and Control* 95 (2018), 1–18.

Loutskina, E. and P. E. Strahan, "Financial Integration, Housing, and Economic Volatility," *Journal of Financial Economics* 115(1) (2015), 1449–1496.

Mian, A., K. Rao and A. Sufi, "House Balance Sheets, Consumption, and the Economic Slump," *Quarterly Journal of Economics* 128 (2013), 1687–1726.

Mian, A. and A. Sufi, "House Prices, Home Equity-Based Borrowing, and the US Household Leverage Crisis," *American Economic Review* 101 (2011), 2132–2156.

———, "House Price Gains and U.S. Household Spending from 2002 to 2006," NBER Working Paper No. 20152, 2014.

Saiz, A., "The Geographic Determinants of Housing Supply," *The Quarterly Journal of Economics* (2010), 1253–1296.

Smets, F. and R. Wouters, "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach," *American Economic Review* 97(3) (2007), 586–606.

Smith, M., "What Do We Know about Equity Withdrawal by Households in New Zealand?," *The Blackwell Companion to the Economics of Housing: The Housing Wealth of Nations* (2010), 176–200.

Sommer, K., P. Sullivan and R. Verbrugge, "The Equilibrium Effect of Fundamentals on House Prices and Rents," *Journal of Monetary Economics* 60(7) (2013), 854–870.

Tauchen, G., "Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions," *Economics Letters* 20 (1986), 177–181.

Wong, A., "Refinancing and the Transmission of Monetary Policy to Consumption," , 2019.

Wooldridge, J. M., "*Econometric Analysis of Cross Section and Panel Data*," Cambridge, Mass: MIT Press, 2002.

Figure 1: Mortgage borrowing, consumption and residential investment (aggregate data)

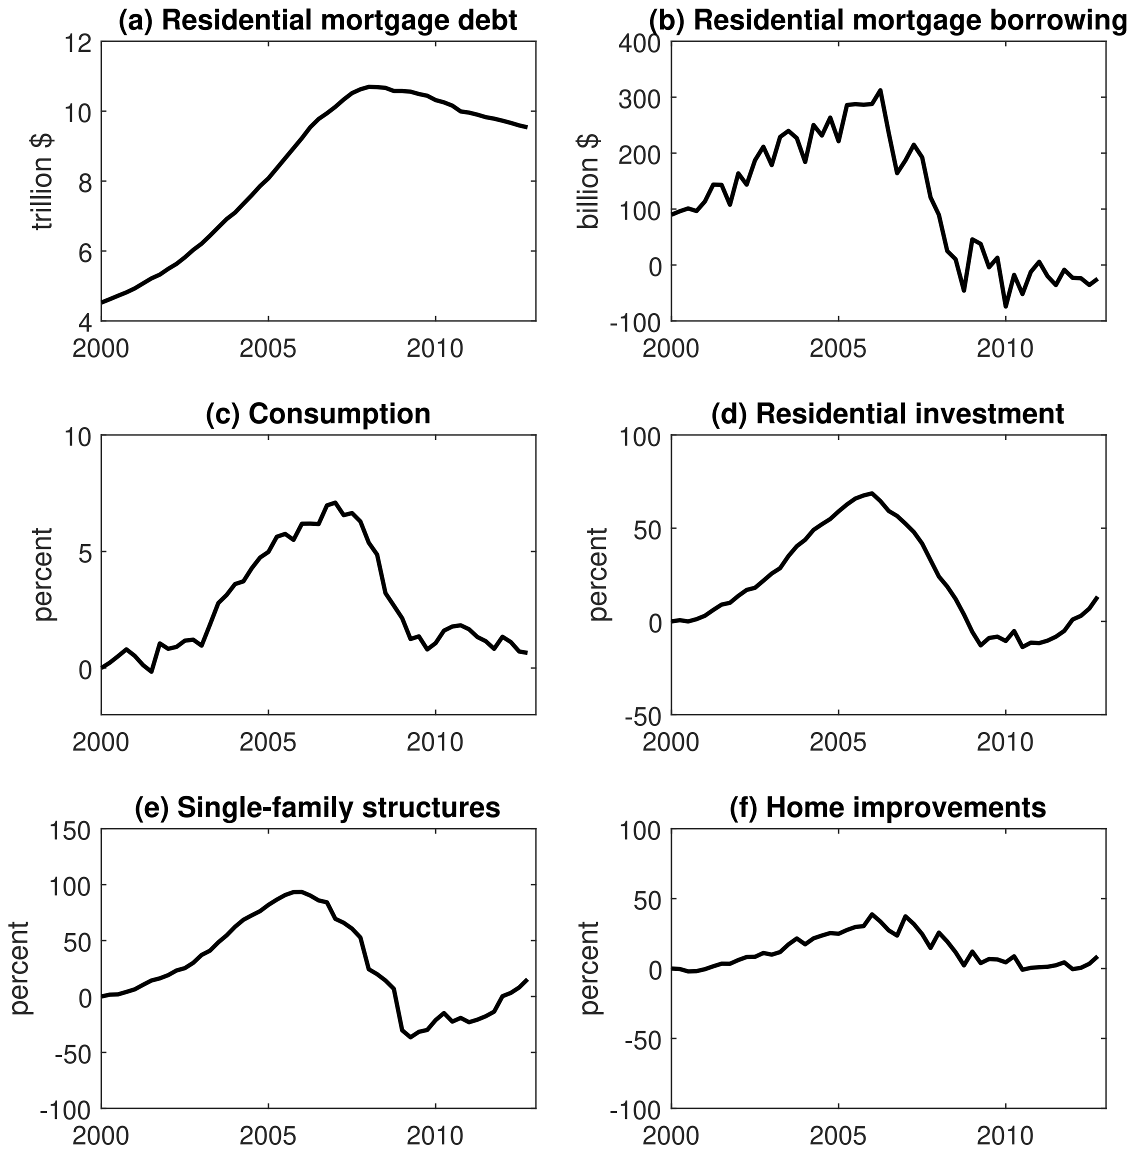

Notes: Residential mortgage debt level and residential mortgage borrowing (converted to quarterly rate) are obtained from Federal Reserve Board Z.1 data release. Data on consumption, residential investment, single-family structures and home improvements are obtained from the Bureau of Economic Analysis. They are deflated by the CPI and are log-linearly detrended. The detrended series are normalized by their corresponding value in 2000.

Figure 2: Life-cycle profiles: Model and microdata

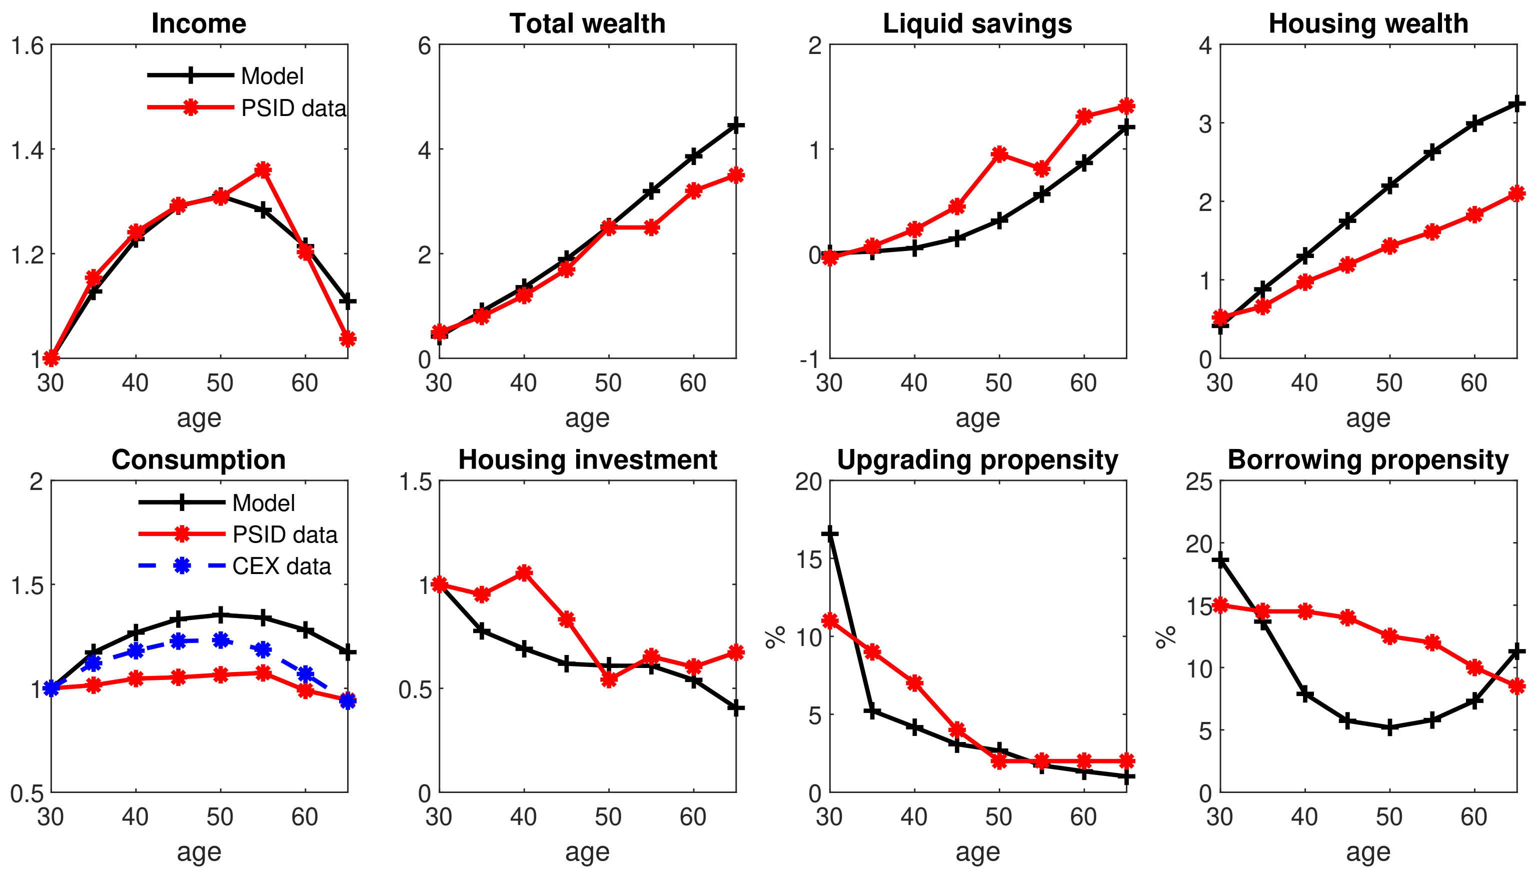

Notes: Life-cycle profiles from household survey data and simulated data from the calibrated model. The income of the youngest age group (25-30) is normalized to 1. Total wealth and its components are in proportion to the income of the youngest age group. Housing investment and consumption are normalized by the corresponding expenditures of the youngest age group. For comparison, a consumption profile is also constructed using the Consumer Expenditure Survey 1994-2014, following the same selection criteria as for the PSID data.

Figure 3: Life-cycle profiles of selected individuals

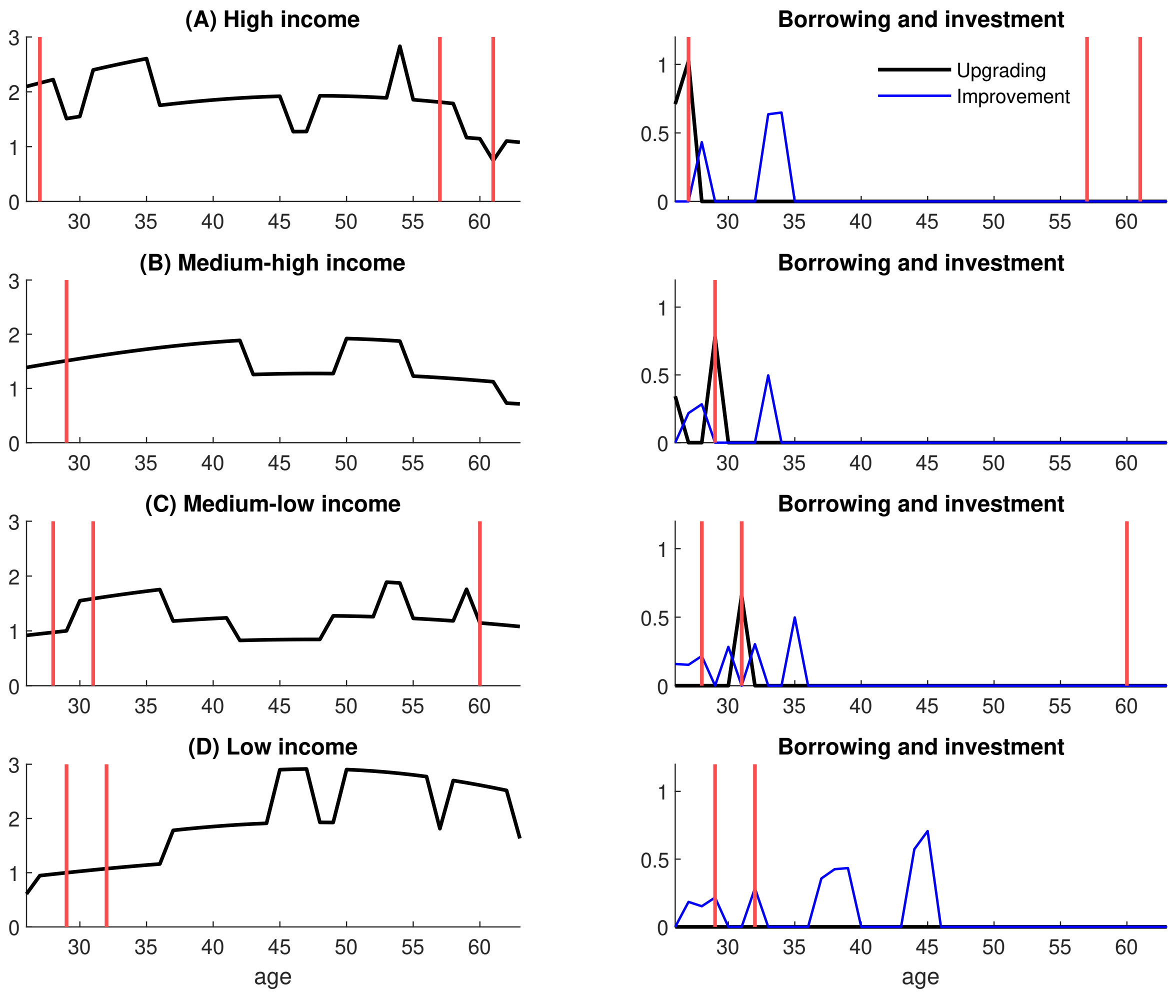

Notes: Simulations from the calibrated model. Red vertical lines indicate the timing of mortgage debt increases.

Figure 4: Responses to 1% increase in house price (model simulation)

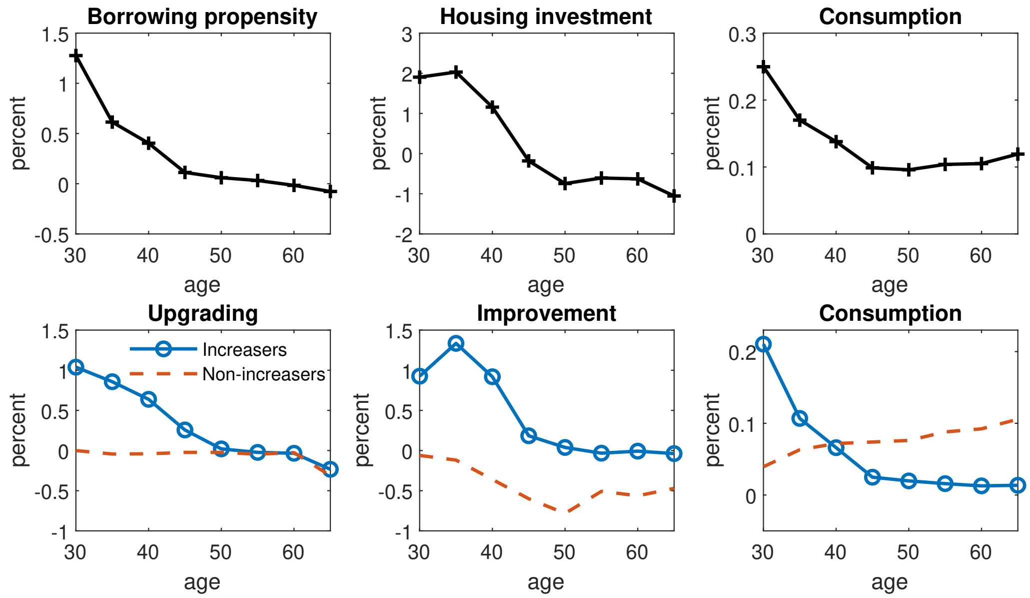

Figure 5: Responses to 1 pp reduction in mortgage rate (model simulation)

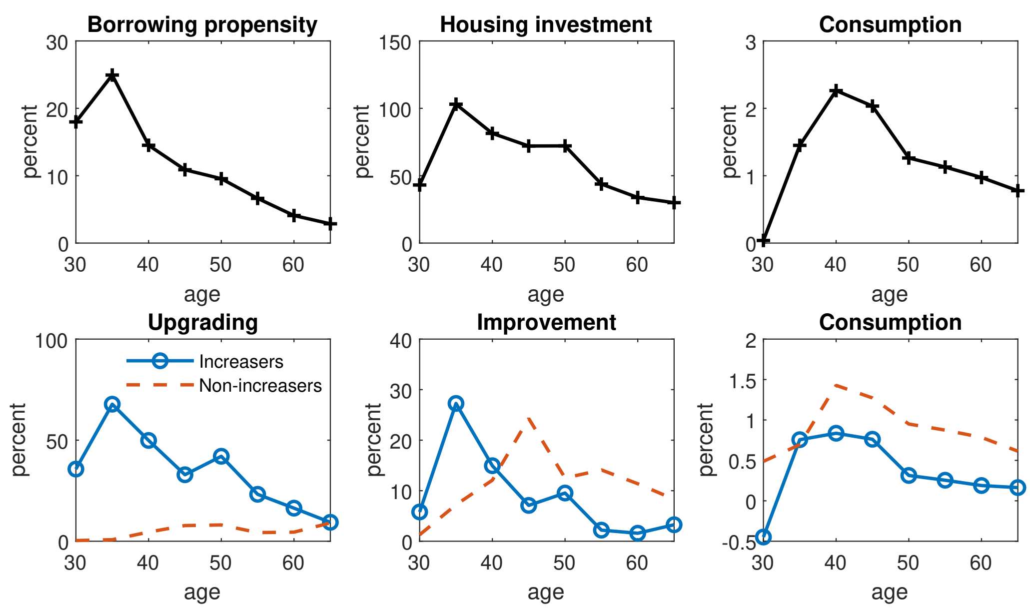

Figure 6: Responses to 1% increase in collateral value (model simulation)

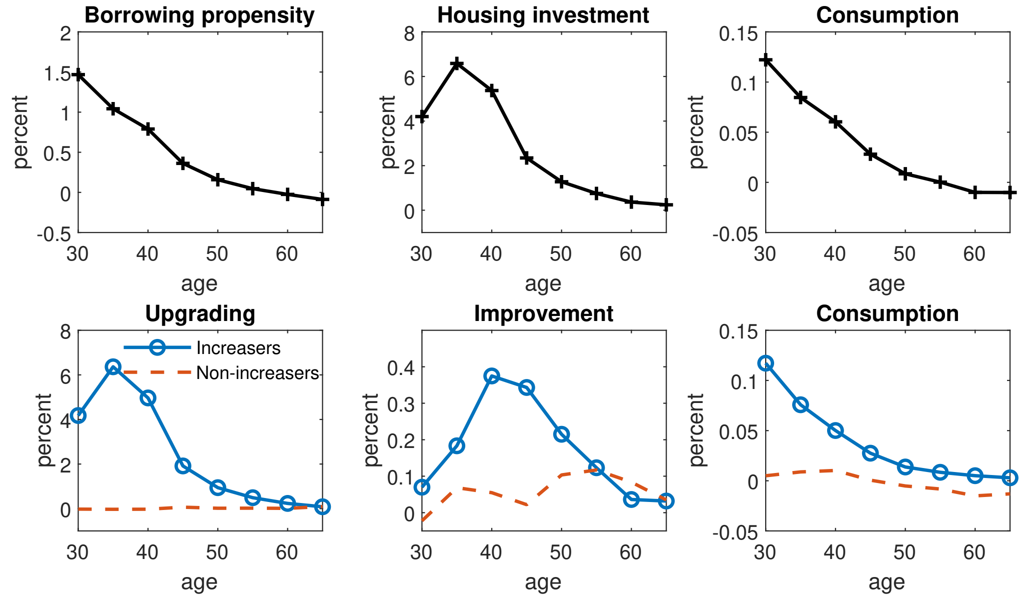

Figure 7: Boom-bust cycles in consumption and residential investment

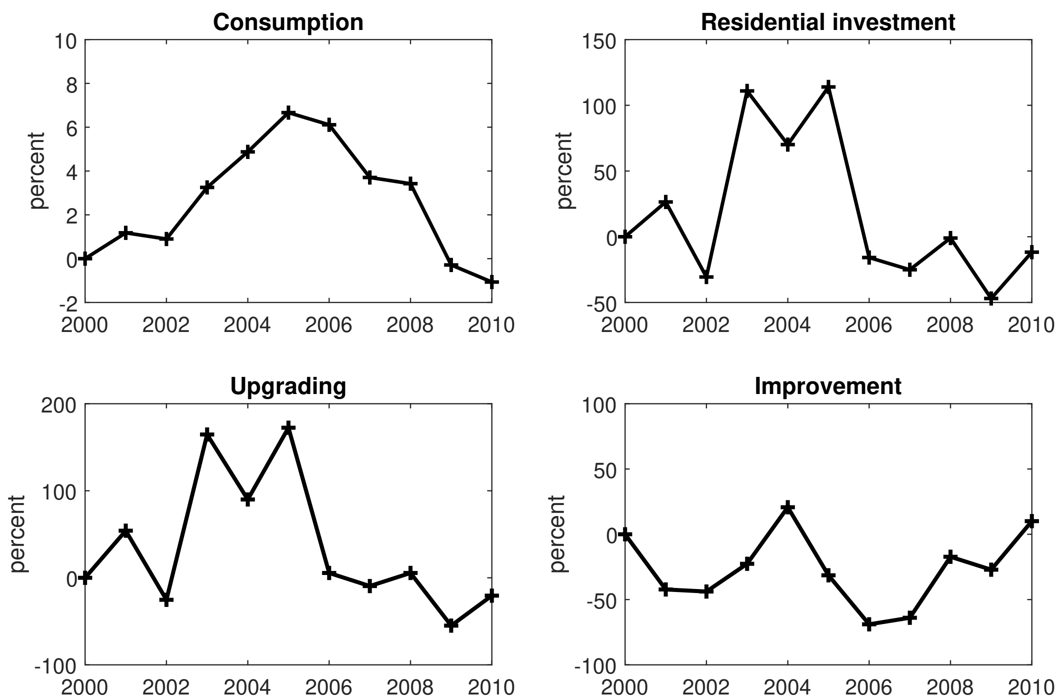

Notes: Simulations from the calibrated model after feeding the historical time series of real house prices, real mortgage rates and a measure of credit conditions into the model. The units are expressed in percent deviations from the corresponding values in 2000.

Figure 8: Expenditure shares of households who increase their mortgage debt

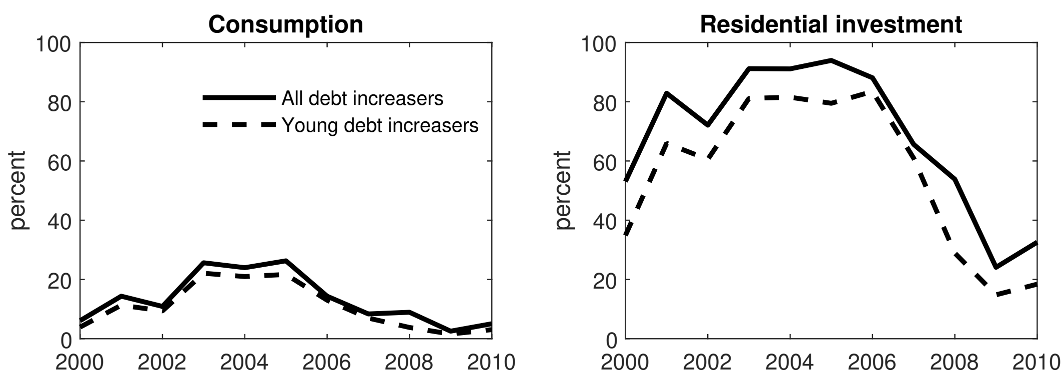

Notes: Simulations from the calibrated model after feeding the historical time series of real house prices, real mortgage rates and a measure of credit condition into the model.

Figure 9: Collateral constraint and aggregate volatility

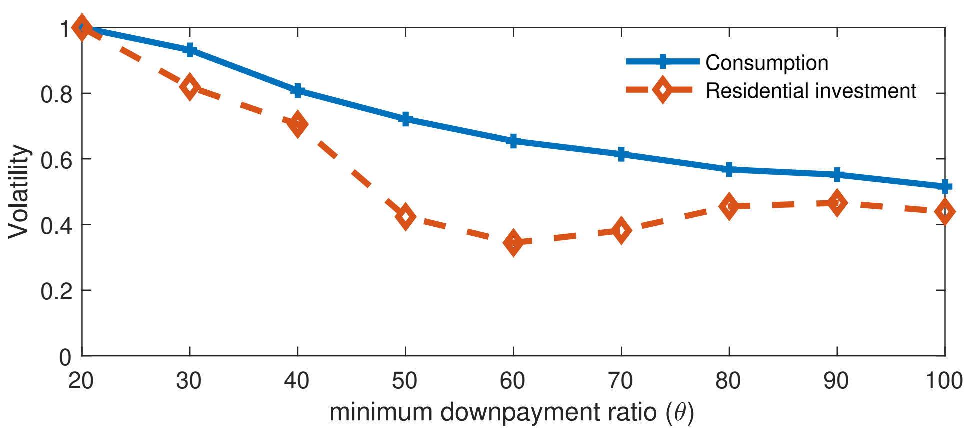

Notes: Simulations from the calibrated model after feeding the historical time series of real house prices, real mortgage rates and a measure of credit condition into the model and setting different values for θ. Volatility is measured by the standard deviation of annual expenditure growth. The volatility measures are normalized to 1 for the actual scenario $(\theta = 0.2)$.

Table 1: Effects of house price shocks: MSA-level evidence

	OLS		2SLS: Saiz instrument		2SLS: GMNS instrument	
	Consumption	Residential investment	Consumption	Residential investment	Consumption	Residential investment
HP growth	0.087***	1.234***	0.143***	0.778***	0.070***	1.031***
	(0.008)	(0.100)	(0.018)	(0.220)	(0.014)	(0.162)
MSA FE	Y	Y	Y	Y	Y	Y
Quarter FE	Y	Y	Y	Y	N	N
Division-by-quarter FE	N	N	N	N	Y	Y
R^2	0.32	0.22	0.39	0.24	0.42	0.27
# Obs.	19,988	19,988	14,159	14,159	16,550	16,550
MPC out of $1 housing wealth (cents)	4.50	5.34	7.40	3.36	3.62	4.46
First-stage						
Saiz elasticity × national HP growth			-0.197***	-0.197***		
			(0.027)	(0.027)		
GMNS sensitivity × national HP growth					1.000***	1.000***
					(0.037)	(0.037)
R^2			0.57	0.57	0.82	0.82
# Obs.			14,159	14,159	16,550	16,550

Notes: *, ** and *** denote significance at the 10%, 5% and 1% level, respectively. Standard errors are clustered at the MSA level. Saiz elasticity refers to the housing supply elasticity estimated by Saiz (2010). GMNS sensitivity refers to the local house price sensitivity estimate in Guren et al. (2020).

Table 2: Effects of mortgage rate shocks : MSA-level evidence

	Consumption	Residential investment	Voluntary prepayment
Panel I. Shock measured by $frac > 0$ (OLS)			
Fraction of loans with positive rate gaps	0.006	0.476***	0.226***
	(0.009)	(0.088)	(0.017)
MSA FE	Y	Y	Y
Quarter FE	Y	Y	Y
R^2	0.24	0.21	0.83
# Obs.	20,110	20,110	20,110
Panel II. Shock measured by average gap (2SLS)			
Average rate gap	0.028**	0.618***	0.296***
	(0.013)	(0.138)	(0.033)
Average rate gap × Post 2008	0.012***	-0.134***	-0.062***
	(0.004)	(0.039)	(0.010)
MSA FE	Y	Y	Y
Quarter FE	Y	Y	Y
R^2	0.23	0.20	0.77
# Obs.	20,110	20,110	20,110

Notes: *, ** and *** denote significance at the 10%, 5% and 1% level, respectively. Standard errors are clustered at the MSA level. The source of the underlying mortgage loan-level data is Black Knight McDash Data.

Table 3: Effects of credit supply shocks : MSA-level evidence

	OLS		2SLS: LS instrument	
	Consumption	Residential investment	Consumption	Residential investment
Credit growth	0.008	0.510***	0.023	0.876***
	(0.005)	(0.070)	(0.050)	(0.307)
MSA FE	Y	Y	Y	Y
Year FE	Y	Y	Y	Y
R^2	0.14	0.25	0.14	0.21
# Obs.	3,129	3,487	3,129	3,487
First-stage				
Around Cutoff$_{m,t-1}$			-3.280***	-4.619***
			(1.048)	(0.971)
Around Cutoff$_{m,t-1}\times$ Change in Cutoff$_t$			0.030	0.119**
			(0.049)	(0.047)
R^2			0.81	0.81
# Obs.			3,129	3,487

Notes: *, ** and *** denote significance at the 10%, 5% and 1% level, respectively. Standard errors are clustered at the MSA level. LS instrument refers to instruments used by Loutskina and Strahan (2015).

Table 4: Implied contributions to GDP growth

	House price shock	Mortgage rate shock	Credit shock
1 s.d. shock in the data	7%	67 bps	27%
Consumption			
Response to a 1 s.d. shock (pp)	0.49	2.28	0.62
Implied contribution to GDP growth	0.32	1.48	0.40
Residential investment			
Response to a 1 s.d. shock (pp)	7.22	36.92	23.65
Implied contribution to GDP growth	0.36	1.85	1.18

Notes: The response to a 1 standard deviation (s.d.) shock is computed by multiplying the preferred estimate for a shock (Tables 1-3) by the shock size (1 s.d.). The implied contribution to GDP growth is computed by multiplying the response to a 1 s.d. shock by the corresponding share in GDP, which is 0.65 for consumption and 0.05 for residential investment.

Table 5: Calibration

Parameter		Value
Preferences		
σ	Inverse of IES	2
β	Discount factor	0.935
α	Consumption expenditure share	0.81
η	Bequest parameter	4
Income		
χ_j	Deterministic component of income	$\xi_1 j + \xi_2 j^2$
ξ_1	Slope parameter of age	0.067
ξ_2	Slope parameter of age-squared	-0.0007
ρ_z	Persistence of idiosyncratic income shocks (z)	0.9
σ_{ez}	Standard deviation of z's error term	0.18
ζ	Retirement income as a fraction of the income in the final working year	0.6
Mortgages		
θ	Minimum down-payment ratio	0.2
f^b	Mortgage refinancing cost	0.02
Housing		
δ	Housing depreciation rate	0.0227
f^{up}	Transaction cost of home upgrading	0.02
f^{imp}	Transaction cost of home improvement	0.01
f^{down}	Transaction cost of downsizing	0.03
$\underline{h}^{up}$	Threshold of incurring home-upgrading transaction cost	0.15
$\underline{h}^{imp}$	Threshold of incurring home-improvement transaction cost	0.07
Prices		
p	Steady state house price	1
λ	Steady state credit condition	1
r^b	Steady state mortgage rate	0.04
r^a	Interest rate of liquid savings	0.01

Table 6: Change in housing investment likelihood associated with mortgage debt increases

	All homeowners	Age group			
		26-35	36-45	46-55	56-65
Panel I. Any housing investment					
Microdata	0.093***	0.159***	0.099***	0.067***	0.070***
	(0.005)	(0.013)	(0.010)	(0.009)	(0.012)
# Obs.	17,097	2,658	4,586	5,587	4,266
Model	0.195***	0.315***	0.219***	0.137***	0.068***
	(0.001)	(0.003)	(0.004)	(0.004)	(0.002)
# Obs.	140,400	32,400	36,000	36,000	36,000
Panel II. Home upgrading					
Microdata	0.064***	0.138***	0.087***	0.039***	0.025***
	(0.003)	(0.012)	(0.007)	(0.005)	(0.004)
# Obs.	17,097	2,658	4,586	5,587	4,266
Model	0.094***	0.203***	0.098***	0.054***	0.031***
	(0.001)	(0.004)	(0.003)	(0.002)	(0.001)
# Obs.	140,400	32,400	36,000	36,000	36,000
Panel II. Home improvement					
Microdata	0.042***	0.059***	0.037***	0.034***	0.053***
	(0.005)	(0.013)	(0.010)	(0.009)	(0.012)
# Obs.	17,097	2,658	4,586	5,587	4,266
Model	0.035***	0.096***	-0.036***	-0.004	0.023***
	(0.003)	(0.006)	(0.008)	(0.007)	(0.002)
# Obs.	140,400	32,400	36,000	36,000	36,000

Notes: Estimates of β_1 in equation (12) using the PSID data and simulated data from the model. *, ** and *** denote significance at the 10%, 5% and 1% level, respectively.

Table 7: Growth in housing investment expenditures associated with mortgage debt increases

	All homeowners	Age group			
		26-35	36-45	46-55	56-65
Panel I. Any housing investment					
Microdata	0.186***	0.382***	0.208***	0.127***	0.068
	(0.016)	(0.051)	(0.036)	(0.019)	(0.048)
# Obs.	16,050	2,537	4,338	5,227	3,948
Model	0.158***	0.225***	0.178***	0.142***	0.098***
	(0.002)	(0.003)	(0.005)	(0.007)	(0.005)
# Obs.	140,400	32,400	36,000	36,000	36,000
Panel II. Home upgrading					
Microdata	0.159***	0.343***	0.171***	0.102***	0.053
	(0.014)	(0.049)	(0.029)	(0.016)	(0.045)
# Obs.	16,050	2,537	4,338	5,227	3,948
Model	0.149***	0.198***	0.195***	0.139***	0.084***
	(0.002)	(0.003)	(0.006)	(0.007)	(0.005)
# Obs.	140,400	32,400	36,000	36,000	36,000
Panel III. Home improvement					
Microdata	0.026***	0.039***	0.036*	0.025***	0.015
	(0.006)	(0.013)	(0.021)	(0.008)	(0.012)
# Obs.	16,050	2,537	4,338	5,227	3,948
Model	0.009***	0.027***	-0.017***	0.003	0.014***
	(0.001)	(0.002)	(0.003)	(0.002)	(0.002)
# Obs.	140,400	32,400	36,000	36,000	36,000

Notes: Estimates of β_1 in equation (13) using the PSID data and simulated data from the model. *, ** and *** denote significance at the 10%, 5% and 1% level, respectively.

Table 8: Housing investment expenditures associated with $1 increase in mortgage debt

	All homeowners	Age group			
		26-35	36-45	46-55	56-65
Panel I. Any housing investment					
Microdata	0.324***	0.672***	0.524***	0.223***	0.045
	(0.042)	(0.095)	(0.092)	(0.045)	(0.138)
# Obs.	16,050	2,537	4,338	5,227	3,948
Model	0.386***	0.815***	0.396***	0.351***	0.215***
	(0.007)	(0.015)	(0.014)	(0.015)	(0.011)
# Obs.	140,400	32,400	36,000	36,000	36,000
Panel II. Home upgrading					
Microdata	0.306***	0.643***	0.474***	0.189***	0.088
	(0.040)	(0.093)	(0.092)	(0.043)	(0.125)
# Obs.	16,050	2,537	4,338	5,227	3,948
Model	0.360***	0.775***	0.394***	0.310***	0.201***
	(0.007)	(0.015)	(0.014)	(0.015)	(0.011)
# Obs.	140,400	32,400	36,000	36,000	36,000
Panel III. Home improvement					
Microdata	0.018	0.029***	0.050*	0.034**	-0.043
	(0.014)	(0.009)	(0.028)	(0.015)	(0.048)
# Obs.	16,050	2,537	4,338	5,227	3,948
Model	0.026***	0.040***	0.002	0.040***	0.014***
	(0.002)	(0.009)	(0.005)	(0.004)	(0.002)
# Obs.	140,400	32,400	36,000	36,000	36,000

Notes: Estimates of β_1 in equation (14) using the PSID data and simulated data from the model. *, ** and *** denote significance at the 10%, 5% and 1% level, respectively.

Table 9: Change in housing investment likelihood associated with mortgage debt increases by income

	All homeowners	High-income	Low-income
Panel I. Home upgrading			
Microdata	0.064***	0.077***	0.050***
	(0.003)	(0.005)	(0.004)
Model	0.094***	0.178***	0.020***
	(0.001)	(0.003)	(0.001)
Panel II. Home upgrading			
Microdata	0.042***	0.036***	0.049***
	(0.005)	(0.008)	(0.007)
Model	0.035***	0.016***	0.059***
	(0.003)	(0.005)	(0.002)

Notes: Estimates of β_1 in equation (12) using the PSID data and simulated data from the model. The high- (low-) income group refers to households whose income growth is above (below) the median of the income growth distribution.

Online Appendix

A Robustness of MSA-level evidence

Table A1: Robustness checks and sub-sample estimates

Panel I. House price shocks					
	Unit measure	p.c. measure	Alt. HPI	1999-2007	2008-2015
Consumption		0.070***	0.079***	0.069***	0.071***
		(0.014)	(0.015)	(0.020)	(0.024)
Residential investment	1.058***	0.967***	0.707***	2.321***	0.931***
	(0.169)	(0.161)	(0.173)	(0.310)	(0.217)
Panel II. Mortgage rate shocks					
	Unit measure	p.c. measure	add. controls	1999-2007	2008-2015
Consumption		0.006	0.001	0.007	0.036***
		(0.009)	(0.008)	(0.014)	(0.013)
Residential investment	0.506***	0.475***	0.470***	0.558***	0.501***
	(0.085)	(0.088)	(0.088)	0.100)	(0.188)
Panel III. Credit supply shocks					
	Unit measure	p.c. measure	add. controls	1999-2006	
Consumption		0.023	0.027	0.009	
		(0.050)	(0.045)	(0.045)	
Residential investment	0.555*	1.074***	1.642***	1.658***	
	(0.308)	(0.342)	(0.306)	(0.268)	

Notes: House price shocks (panel I) are identified using the GMNS instrumental-variable approach. Mortgage rate shocks (panel II) are measured by $frac > 0$. Credit shocks (panel III) are identified using the LS instruments. Column (1) shows the responses of residential investment measured by annual growth of building permit units. Column (2) shows the responses of residential investment measured by annual growth of building permit valuation per capita. In column (3) of panel I, the FHFA house price indices are used. In column (3) of panel II, additional controls are included (the interaction between Guren et al.'s sensitivity measure and national house price growth, and the differential sensitivity of retail employment to regional retail employment and to Gilchrist and Zakrajsek (2012) excess bond premia). In column (3) of panel III, additional controls are included (the interaction between Guren et al.'s sensitivity measure and national house price growth and the differential sensitivity of retail employment to regional retail employment and to the 30-year fixed mortgage rate). The last two columns present sub-sample estimates from the preferred specifications.

B PSID Data and Additional Empirical Results

The PSID data for 1999-2015 are used to provide household-level evidence in Section 6.1 and household-level evidence for the effects of house price shocks in Section 3.1. The data are representative of the U.S. population and contain detailed information on wealth, income, and expenditures of a panel of approximately 7,000 households, making it more suitable than other U.S. household surveys to study consumer borrowing and spending patterns.[27]

I applied the following sample selection criteria. First, I focus on households of working age (i.e., the heads of the household between age 26 and 65). Second, I focus on existing homeowners, because they have the ability to use their homes as collateral to borrow.[28] I also exclude first-time home buyers, because the increase in their mortgage debt (from zero) is mechanically used to finance their housing investment. Including them in the sample will reinforce my results. Third, I exclude households owning farms or businesses, because their borrowing may be related to their business investment, rather than consumer spending. Fourth, I exclude households having negative total income or a home value below $5,000 to reduce possible errors in the survey data. Below I describe the definition of the key variables in the empirical analysis. See Table B1 for the summary statistics of selected key variables.

- Indicator for increasing one's mortgage debt: A homeowner's total mortgage balance increases by at least 5% between two survey periods, and the increase exceeds $1,000 (as used in Bhutta and Keys (2016)). Note that total mortgage balance refers to the first and second liens on the primary residence, as the PSID does not collect information on mortgages secured by non-primary residence.
- Indicator for home upgrading: A homeowner sells the old home and moves to a more expensive new home within the same period. For robustness, I also use alternative definitions based on the number of main rooms (see Table B4).
- Expenditures on home upgrading: The difference between the market price of the current

[27] The Survey of Consumer Finance (SCF), for example, collects data on household wealth but not on expenditures. In addition, SCF data are cross-sectional with the exception of the 2007-09 panel that was designed to study the impact of the Great Recession. The Consumer Expenditure Survey (CE) collects detailed expenditure data over a four-quarter panel but has limited information on wealth and income. In particular, it does not reflect the change in outstanding mortgage balances, making it impossible to distinguish, for example, between cash-out refinancing and rate refinancing. Moreover, it stops tracking households who move to another place during the interview period.

[28] This selection is also motivated by the empirical observation in Mian and Sufi (2011) and Bhutta (2015) that household debt dynamics in the 2000s were mainly driven by existing homeowners.

home and the price at which the old home was sold (all in 2009 dollars). This definition is used to match the concept of housing investment in the theoretical model, i.e., $p_t \left[h_{i,t} - h_{i,t-1}(1-\delta)\right]$.

- Indicator for home improvement: A homeowner makes any additions, remodeling or improvements that exceed $10,000 (not including general maintenance or upkeep).[29]
- Expenditures on home improvement: Expenditures reported by the homeowner making home improvements.
- Consumption expenditures include: (i) non-durable goods: food (at home, delivered, and eaten out), gasoline, and clothing (available since 2005); (ii) durable goods: automobiles, furniture (available since 2005), and trips and recreation (available since 2005); (iii) services: utilities (heating, electricity, water, and other kinds of utilities), education, health (hospital bills, doctor visits, prescription drugs, and health insurance), child care, vehicle lease (lease payment and fees), public transportation (bus, train, taxi, and other transportation services), parking, vehicle repair, vehicle insurance, homeowner insurance, property tax, telephone (available since 2005), internet (available since 2005), and home repair services (available since 2005).
- Total household income: taxable income, transfer income and social security income of the head of household and of the spouse in the previous year.
- Wealth: (i) net liquid assets: sum of safe assets (checking and savings accounts, money market funds, certificates of deposit, government savings bonds, and treasury bills) and risky assets (stock shares) net of non-mortgage debt (credit cards, student loans, medical or legal bills, and loans from relatives); (ii) net illiquid assets: sum of home equity, net value of real estate owned by the household, private annuities and individual retirement accounts, net value of other assets, and the value of farms or businesses. Note that this definition is used when wealth is included as a control variable in a regression. When comparing the calibrated model with the data (Figure 2), wealth is defined as the sum of liquid assets and home equity.

[29] In the PSID survey, home improvement expenditures are reported when exceeding $10,000, which leads to a truncated-variable problem. Since I estimate the change in expenditure growth associated with debt increases, this problem cannot be addressed by estimating the Tobit model. Note that, since the 2005 survey, home maintenance expenditures (often less than $10,000) are reported without truncation. I conduct robustness checks that impute the home maintenance expenditures of households prior to the 2005 survey by matching their characteristics to households after the 2005 survey. I then redefine home improvements as the sum of the reported improvements and the imputed maintenance expenditures. The estimates based on this measure exhibit patterns similar to the baseline estimation.

Table B1: Summary statistics, PSID 1999-2015

Year	# Households	Home value ($,000)	Mortgage balance ($,000)	Deb increasers (% of h.h.)	Amount increased ($,000)	Housing investment ($,000) non-increasers	Housing investment ($,000) Increasers	Consumption ($,000) non-increasers	Consumption ($,000) Increasers
		(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
99-01	1,890	146.5	85.3	27.7	40.0	4.6	22.8	82.9	84.4
01-03	2,022	170.5	99.2	33.4	40.7	5.1	23.4	81.0	93.6
03-05	2,148	219.1	119.0	32.3	54.4	4.4	34.8	84.8	92.6
05-07	2,167	248.5	137.5	31.4	64.7	6.4	35.4	83.2	89.6
07-09	2,215	217.5	141.6	23.5	51.3	4.1	15.7	81.0	88.8
09-11	2,250	205.6	145.2	19.1	49.0	2.6	15.1	79.7	82.8
11-13	2,254	207.4	146.2	15.6	49.2	2.7	14.5	71.0	75.5
13-15	2,151	224.5	145.1	16.4	61.7	3.6	28.3	72.5	78.9
Average		206.2	128.4	24.7	51.4	4.1	25.0	79.2	87.4

Notes: Statistics in columns (1) to (8) are reported as household average. Mortgage balances in (2) are conditional on households having mortgages. The amount borrowed in (4) is conditional on households who increased their mortgage debt. Housing investment and consumption expenditures in (5)-(8) are converted to 2009 dollars.

Table B2: Effects of house price shocks: Evidence from PSID data

	All homeowners	Age group			
		26-35	36-45	46-55	56-65
Panel I. Consumption					
HP growth	0.086***	0.010	0.138**	0.062*	0.113***
	(0.016)	(0.042)	(0.056)	(0.037)	(0.035)
# Obs.	16,050	2,537	4,338	5,227	3,948
Panel II. Likelihood of making housing investment					
HP growth	0.151***	0.239***	0.178***	0.124***	0.082***
	(0.009)	(0.023)	(0.018)	(0.014)	(0.018)
# Obs.	16,050	2,537	4,338	5,227	3,948
By investment type: (i) Home upgrading					
HP growth	0.082***	0.196***	0.103***	0.051***	0.033***
	(0.005)	(0.018)	(0.011)	(0.006)	(0.006)
# Obs.	16,050	2,537	4,338	5,227	3,948
(ii) Home improvement					
HP growth	0.056***	0.040***	0.063***	0.067***	0.046***
	(0.007)	(0.013)	(0.013)	(0.012)	(0.016)
# Obs.	16,050	2,537	4,338	5,227	3,948
By borrowing status: (i) Those increasing mortgage debt					
HP growth	0.218***	0.243***	0.258***	0.209***	0.148***
	(0.019)	(0.039)	(0.034)	(0.037)	(0.037)
# Obs.	3,996	735	1,234	1,282	745
(ii) Those not increasing mortgage debt					
HP growth	0.059***	0.084***	0.066***	0.065***	0.042**
	(0.010)	(0.028)	(0.021)	(0.014)	(0.021)
# Obs.	12,054	1,802	3,104	3,945	3,203
Panel III. Likelihood of increasing mortgage debt					
HP growth	0.263***	0.471***	0.315***	0.182***	0.147***
	(0.013)	(0.033)	(0.029)	(0.020)	(0.020)
# Obs.	16,050	2,537	4,338	5,227	3,948

Notes: Responses to a 1% increase in home value. *, ** and *** denote significance at the 10%, 5% and 1% level, respectively. Standard errors are clustered at the household level. All regressions control for household demographics, financial conditions in the previous period, changes in income and mortgage rates, and household and year fixed effects.

Table B3: Mortgage debt increases and change in expenditures

	Total expenditures	Consumption	Housing investment
Panel I. Expenditure growth associated with debt increases			
Microdata	0.161***	-0.025	0.186***
	(0.029)	(0.025)	(0.016)
# Obs.	16,050	16,050	16,050
Model	0.164***	0.006***	0.158***
	(0.002)	(0.001)	(0.002)
# Obs.	140,400	140,400	140,400
Panel II. Expenditures associated with $1 increase in debt			
Microdata	0.319***	-0.005	0.324***
	(0.045)	(0.015)	(0.042)
# Obs.	16,050	16,050	16,050
Model	0.395***	0.009***	0.386***
	(0.007)	(0.001)	(0.006)
# Obs.	140,400	140,400	140,400

Notes: Panels I and II show the estimates of β_1 in equations (13) and (14), respectively, using the PSID data and simulated data from the model. *, ** and *** denote significance at the 10%, 5% and 1% level, respectively.

Table B4: Alternative measures of home upgrading in PSID

	All homeowners	Age group			
		26-35	36-45	46-55	56-65
Baseline: moving to a more expensive home					
Change in likelihood	0.064***	0.138***	0.087***	0.039***	0.025***
	(0.003)	(0.012)	(0.007)	(0.005)	(0.004)
Change in exp. growth	0.162***	0.339***	0.172***	0.097***	0.077**
	(0.013)	(0.047)	(0.028)	(0.016)	(0.032)
Change in exp. amount	0.307***	0.643***	0.475***	0.190***	0.089
	(0.040)	(0.093)	(0.092)	(0.044)	(0.125)
Alt. 1: moving to a more expensive home with at least same # of main rooms					
Change in likelihood	0.053***	0.128***	0.077***	0.027***	0.017***
	(0.003)	(0.013)	(0.008)	(0.004)	(0.004)
Change in exp. growth	0.110***	0.196***	0.130***	0.052***	0.066*
	(0.014)	(0.056)	(0.029)	(0.014)	(0.039)
Change in exp. amount	0.191***	0.404***	0.319***	0.155***	0.012
	(0.045)	(0.125)	(0.072)	(0.050)	(0.152)
Alt. 2: moving to a home with at least one more main room					
Change in likelihood	0.041***	0.109***	0.061***	0.018***	0.011***
	(0.003)	(0.012)	(0.007)	(0.004)	(0.003)
Change in exp. growth	0.080***	0.152***	0.101***	0.029**	0.037
	(0.014)	(0.052)	(0.035)	(0.013)	(0.037)
Change in exp. amount	0.166***	0.339***	0.287***	0.127***	0.010
	(0.043)	(0.127)	(0.064)	(0.042)	(0.148)
Alt. 3: moving to a more expensive home with at least one more main room					
Change in likelihood	0.042***	0.113***	0.064***	0.018***	0.010***
	(0.003)	(0.013)	(0.008)	(0.004)	(0.003)
Change in exp. growth	0.083***	0.151***	0.104***	0.030***	0.039
	(0.013)	(0.052)	(0.028)	(0.011)	(0.036)
Change in exp. amount	0.156***	0.339***	0.276***	0.108***	0.008
	(0.042)	(0.127)	(0.064)	(0.039)	(0.148)

Notes: Change in likelihood refers to the estimates of β_1 in equation (12). Change in exp. growth refers to the estimates of β_1 in equation (13). Change in exp. amount refers to the estimates of β_1 in equation (14).

C Numerical Solution Methods

I use a two-step numerical procedure to solve the stationary equilibrium of the model. In the first step, I solve the value functions indexed by j over a fixed-grid state space characterized by (h, b, a, y). I choose 20 grids for h, 20 grids for b, and 10 grids for a. For the stochastic component of the logarithm of income, z, I use the Tauchen (1986) method to discretize an AR(1) process to 5 states. The Tauchen method chooses grid points and computes the corresponding transitional probability matrix. The following backward-induction algorithm is used to solve the value functions (step 1):

1. The value function of the last period, V_J, is determined by the bequest function.
2. For each j and a given state (h, b, a, y), I compute $V_j^R(h, b, a, y)$ and $V_j^N(h, b, a, y)$ separately. To obtain $V_j^R(h, b, a, y)$, I construct a choice set $\left\{h_R^{'}(h, b, a, y), b_R^{'}(h, b, a, y), a_R^{'}(h, b, a, y)\right\}$ that takes into account all constraints and may vary across states.
3. I then compute the value of the objective function for each choice. To compute the continuation value, I use the linear interpolation if the choice is not one of the grid points.
4. I set $V_j^R(h, b, a, y)$ to the maximum value of the objective function after evaluating all choices.
5. A similar procedure is used to obtain $V_j^N(h, b, a, y)$. Then, the value function at age j is set to $V_j(h, b, a, y) = \max\{V_j^R(h, b, a, y), V_j^N(h, b, a, y)\}$.

In the second step, given j, V_j (obtained in step 1), and a realized state (h, b, a, y), I solve the optimal choices, $(h_R^{'}, b_R^{'}, a_R^{'})$ and $(h_N^{'}, a_N^{'})$, respectively, over a finer choice set that has four times as many as the choices for each variable, using a similar approach as described in 2-5 above. This procedure allows me to find the solution more accurately.

To obtain the impact responses to a permanent aggregate shock, I first solve the set of value functions at the new value of the aggregate state (step 1). I then set the individual-specific state at the beginning of the shock period equal to the value in the original stationary equilibrium, and solve the optimal choices with the new value of the aggregate state and the new continuation values obtained from step 1.

D Canonical Model with One Type of Housing Investment

This section calibrates a life-cycle model that is identical to the model in Section 4, except that there is no distinction between home upgrading and home improvements. Specifically, households pay a transaction cost if the new home size (after an investment) relative to the current home size exceeds a certain threshold. The transaction cost is set to the middle point of the transaction cost of upgrading and of improvement, i.e., 0.015. The threshold is calibrated to match the annual investment rate (the frequency of upgrading or making an improvement), which is 12%.

Figure D1 shows key life-cycle profiles of the average household simulated from the model compared to household survey data. The model matches well these profiles in the data, with the exception of the investment propensity. The model implies a much higher investment rate for young homeowners and a much lower investment rate for old households. Changing the parameters to dampen the investment rate for the young would result in even lower investment rates for the old, and vice versa. The model in Section 4, in contrast, imposes differential investment parameters for home upgrading, which helps to limit young households' investment frequency.

Figure D2 plots the impact responses of borrowing and spending to each of the shocks in Section 7. These responses are similar to those in Figures 4 to 6, indicating that the model with two investment types captures the predictions in the canonical model about the overall spending and borrowing responses.

To summarize, it is hard to match the investment aspect of the canonical model directly with microdata, because different types of investment do exist, with moving and improvements most commonly seen. Previous studies have used statistics about moving to calibrate the investment parameters, but this approach leaves home improvements unexplained. The model in Section 4 adds realism to the canonical life-cycle framework. It also shows that, although home improvements account for an important share of housing investment and are performed frequently, they are not the main component that responds to shocks (see Section 6.2 for the mechanisms). Moreover, I show that the predictions of my model about borrowing, consumption, and overall housing investment are consistent with the canonical life-cycle model.

Figure D1: Life-cycle profiles: Canonical model and microdata

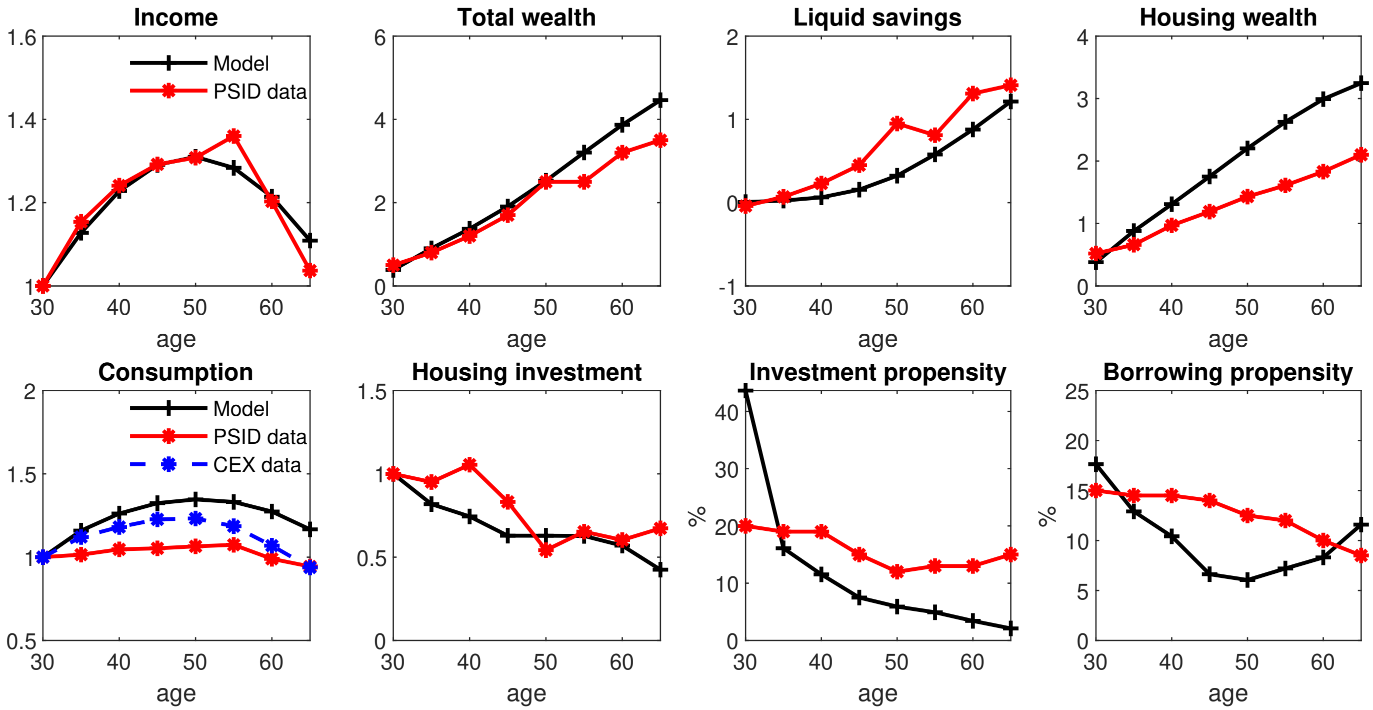

Notes: Simulations are from the model with one investment type. See notes in Figure 2 for detailed variable description.

Figure D2: Impact responses in canonical model

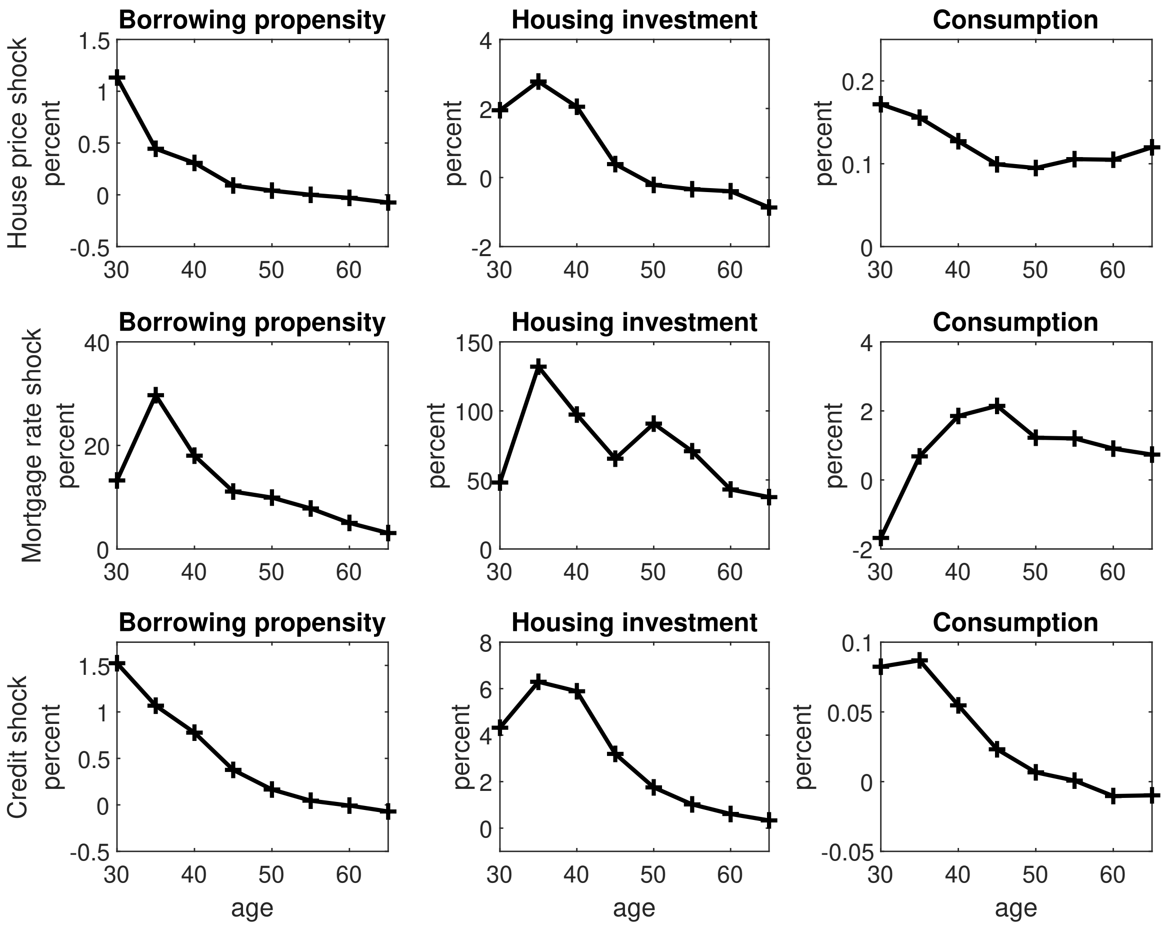

Notes: Simulations from the model with one investment type. The size of each shock is the same as in Figures 4 to 6.

E Impulse Responses: Life-Cycle Model and DSGE Model

This section presents the impulse responses of key aggregate variables from the structural life-cycle model of Section 4 and compares them with the two-agent DSGE model of Iacoviello and Neri (2010) (henceforth the IN model). I choose the IN model for comparison because it embeds a housing sector and constrained households into the mainstream DSGE framework (see Christiano et al. (2005) and Smets and Wouters (2007)). The IN model fits the U.S. data quantitatively well and has become the workhorse DSGE model for studying the interaction between housing and the macroeconomy. The objective is to see whether the structural life-cycle model can generate aggregate responses consistent with DSGE models featuring housing and household heterogeneity.

Both models share a common transmission channel: Shocks to the housing market affect some households' borrowing capacity and therefore their spending. On the other hand, the key difference is that the life-cycle model is partial-equilibrium and does not incorporate (i) the responses of other sectors to a shock, (ii) the redistribution between constrained and unconstrained households, and (iii) the feedback effects of (i) and (ii) on household expenditures. For example, fluctuations in house prices and mortgage rates are exogenous in the life-cycle model but are endogenously determined in the IN model with feedback effects on each other. Despite this difference, however, as I show below, the two models generate qualitatively similar responses. In particular, the responses of young households in the life-cycle model are similar to constrained households in the IN model.

Credit shocks. Figure E1 shows the impulse responses of aggregate borrowing, consumption and residential investment to a 1% credit shock in the IN model (left) and in the life-cycle model (right).[30] This shock has the same interpretation across the two models, making the comparison straightforward. An expansion in credit increases household borrowing, consumption and residential investment in both models. In particular, consumption and residential investment of young/constrained households are more responsive than old/unconstrained households.

Mortgage rate shocks. Since the mortgage rate is endogenously determined in the IN model, to facilitate the comparison, I consider a monetary policy shock in the IN model that lowers the

[30] The quarterly impulse responses from the IN model are averaged to obtain annual responses for the comparison with the life-cycle models. The residential investment responses of borrowers from the IN model are scaled by their labor income share to make the magnitudes of aggregate residential investment visible. The responses of residential investment in the life-cycle model refer to the responses of home upgrading and home improvements, consistent with the impact responses of housing investment in Section 7.

mortgage rate by 1 pp in the first year. Figure E2 shows the impulse responses from the two models. A lower mortgage rate stimulates borrowing, consumption and residential investment in both models, especially for young/constrained households. One difference is that consumption increases more in the IN model, whereas residential investment increases more in the life-cycle model. This is driven by the endogenous response of house prices in the IN model. As the mortgage rate falls, house prices increase, stimulating consumption (due to the collateral effect and the substitution effect), while dampening housing investment to some extent.

House price shocks. Since house prices are endogenously determined in the IN model, to facilitate the comparison, I consider a housing technology (supply) shock in the IN model that raises the real house price by 1% in the first year.[31] Figure E3 shows the impulse responses from the two models. Higher house prices stimulate household borrowing and consumption, especially for young/constrained households. Residential investment falls in the IN model, mainly because unconstrained households cut back on such spending, although constrained households see a rise immediately after the shock. Aggregate housing investment in the life-cycle model shows almost no change, but it increases for young households after the shock, which can be explained by life-cycle demand for housing and the collateral effect.

[31] House price fluctuations can also be driven by housing preference (demand) shocks in the IN model. A positive housing demand shock raises house prices, borrowing, consumption and residential investment, especially for constrained households.

Figure E1: Impulse responses to a credit shock

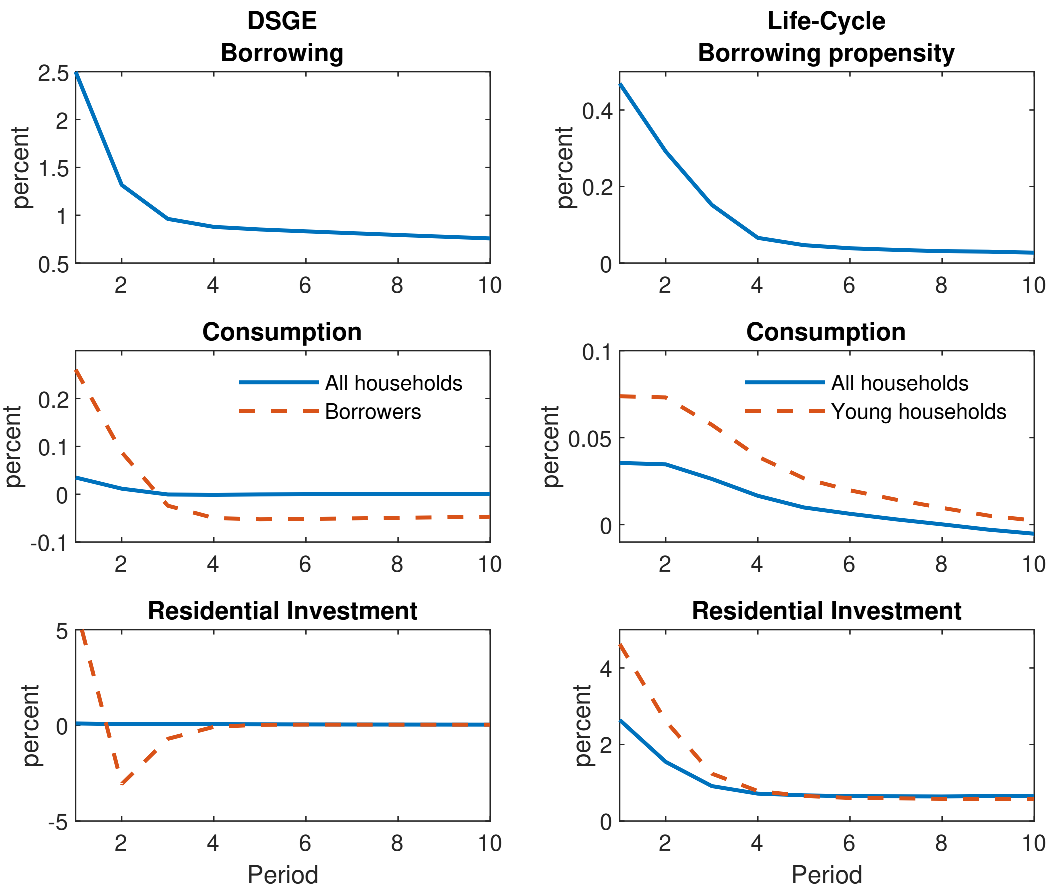

Notes: The left column shows the impulse responses to a 1% credit shock in the DSGE model of Iacoviello and Neri (2010). The right column shows the impulse responses to a 1% credit shock in the life-cycle model of Section 4.

Figure E2: Impulse responses to a mortgage rate shock

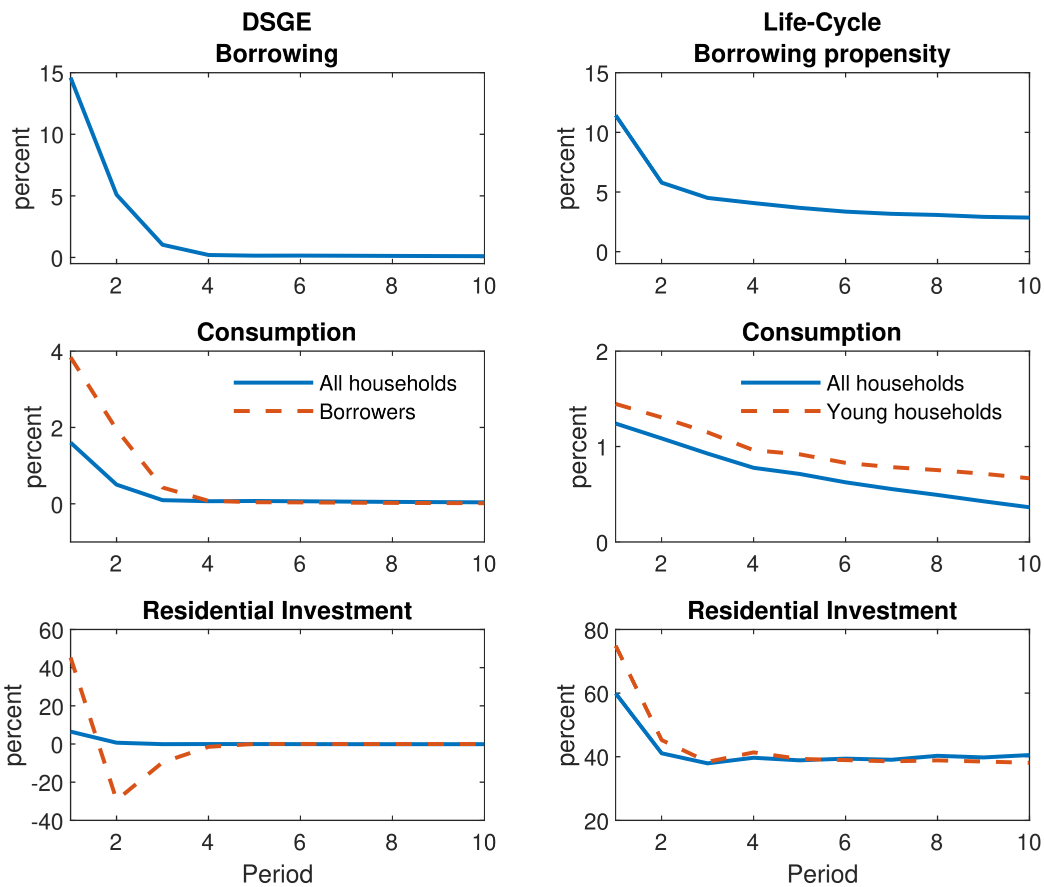

Notes: The left column shows the impulse responses to a monetary policy shock that lowers the mortgage rate by 1 pp in the first year in the DSGE model of Iacoviello and Neri (2010). The right column shows the impulse responses to a 1 pp reduction in the mortgage rate in the life-cycle model of Section 4.

Figure E3: Impulse responses to a house price shock

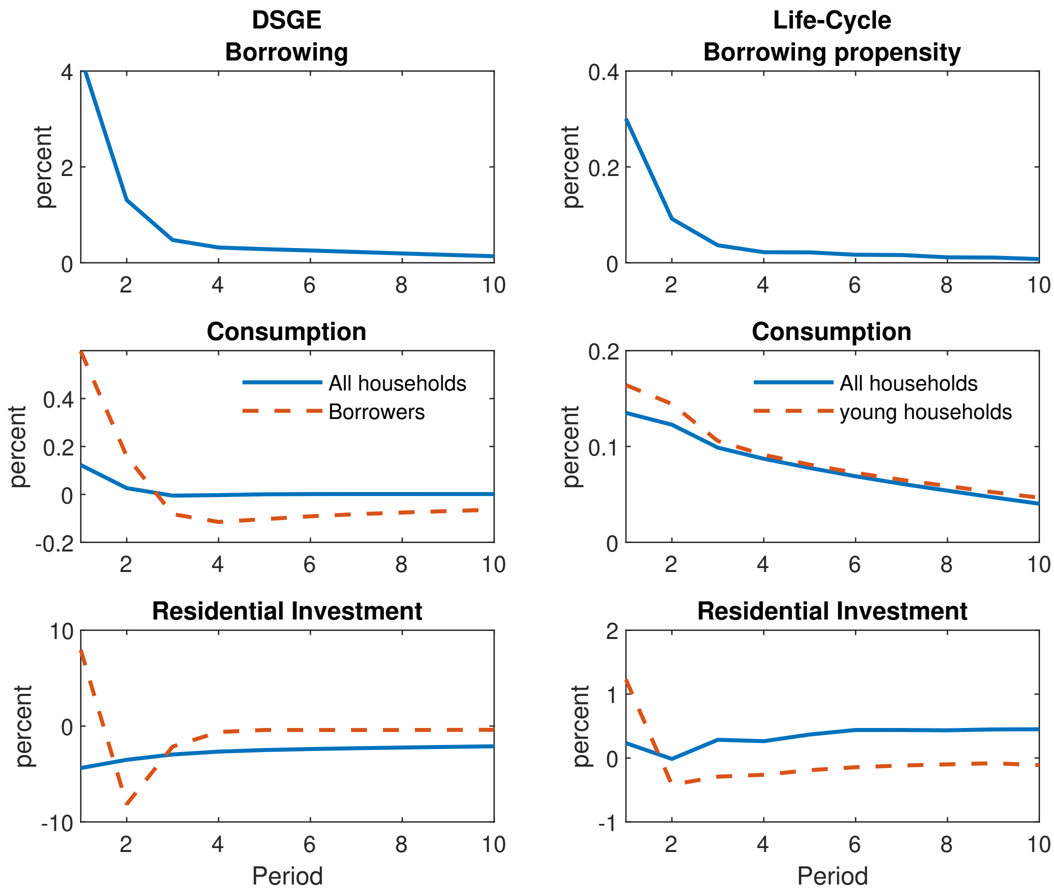

Notes: The left column shows the impulse responses to housing technology shock that raises the real house price by 1% in the first year in the DSGE model of Iacoviello and Neri (2010). The right column shows the impulse responses to a 1% increase in the house price in the life-cycle model of Section 4.

F Accounting for House Price Changes Driven by Demand Shocks

This section complements the analysis in Section 7 by allowing mortgage rate and credit supply shocks to affect house prices. As a consequence, consumption and residential investment respond directly to a change in the mortgage rate or the credit condition, and indirectly through the change in the house price.

To calibrate the change in house prices to these shock, I use MSA-level data and the empirical strategies in Section 3 to estimate the causal responses. Specifically, for the mortgage rate shock, I estimate a 2SLS specification as in question (5) fitting the log annual change in MSA-level house prices. The regression using pre-crisis data gives $\hat{\beta}_1 = 0.16$ (with t-statistics of 6.29). This means that a 1 pp decline in the mortgage rate increases the house price by 16%. For the credit supply shock, I estimate a 2SLS specification as in question (6) with the log annual change in MSA-level house prices as the dependent variable, which gives $\hat{\beta}_1 = 0.14$ (with t-statistics of 1.98), implying that 1% increase in credit raises house prices by 0.14%.

Figure F1 shows the simulated responses to a 1 pp decline in the mortgage rate and a simultaneous increase in house prices of 16%. These responses are similar to Figure 5 with three noticeable differences: (i) a higher borrowing propensity due to increased home equity, especially for young homeowners, (ii) a slightly lower housing investment, especially for middle-age and old homeowners on home improvements, because of higher cost of investing, and (iii) a higher consumption due to increased home equity for all homeowners, especially young homeowners.

Figure F2 shows the simulated responses to a 1% increase in collateral value and a simultaneous increase in house prices of 0.14%. These responses are similar to those in Figure 6 with two noticeable differences: (i) A lower housing investment, especially for middle-age and old homeowners on home improvements, because of higher cost of investing, and (ii) a higher consumption for young homeowners due to increased home equity. Overall, these responses are similar to those presented in Sections 7.2 and 7.3, and the differences reflect the responses to house price changes, as discussed in Section 7.1.

Figure F1: Responses to mortgage rate shocks that also trigger increases in house prices

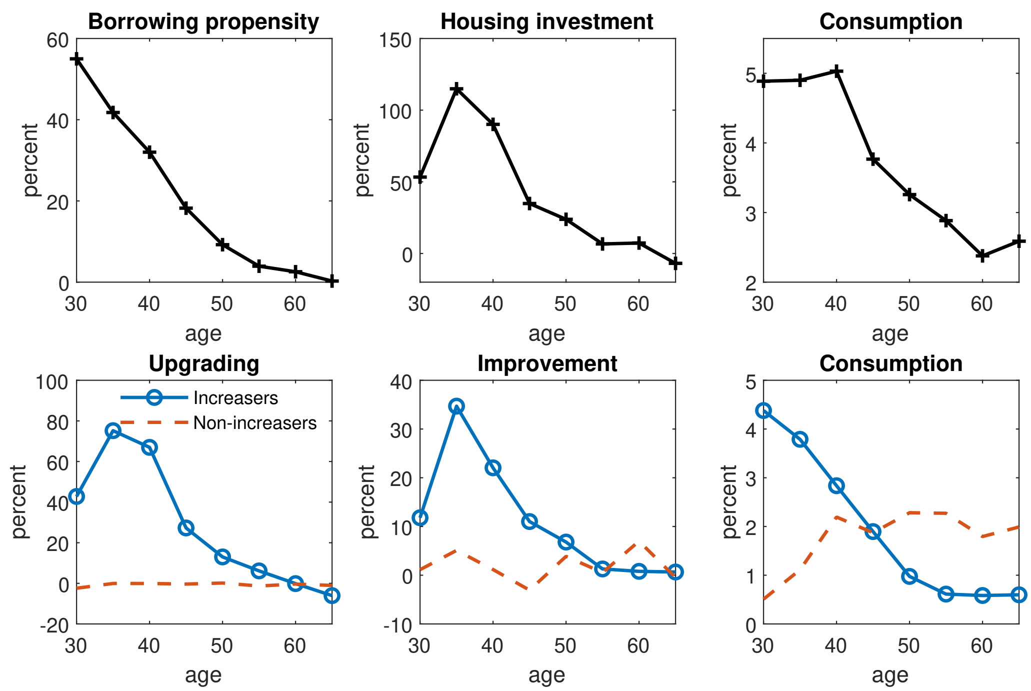

Figure F2: Responses to credit supply shocks that also trigger increases in house prices

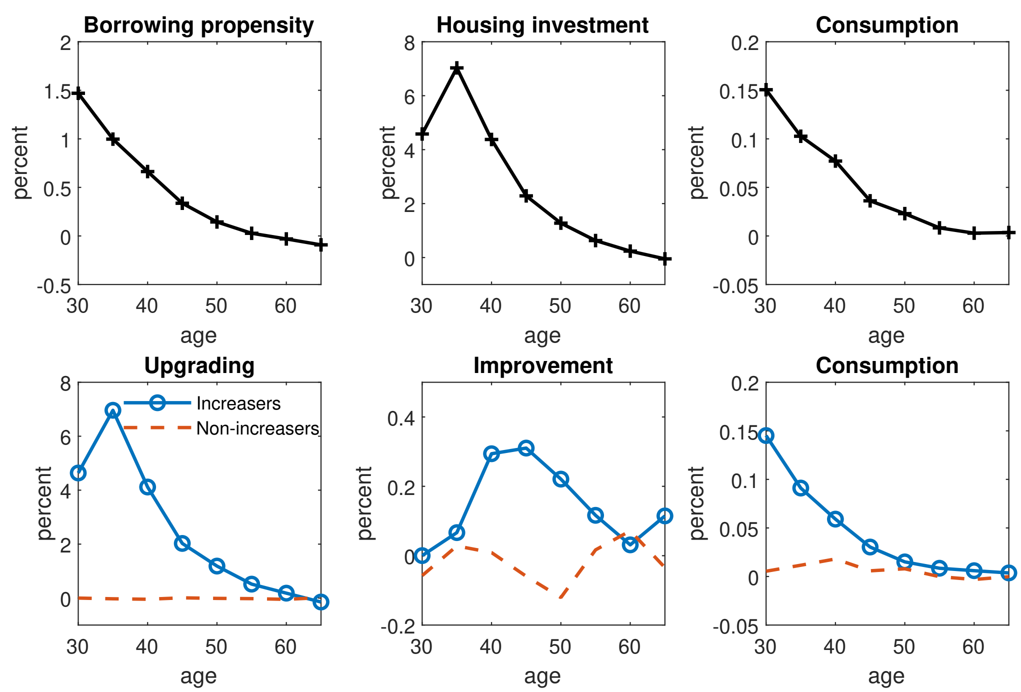

G The Role of Expectations

Recent studies have stressed the role of beliefs about future house prices in driving household spending decisions (see Kaplan et al. (2019) and Bailey et al. (2019)). Intuitively, when households expect future house prices to grow, owning a home becomes more attractive, which stimulates housing investment. This optimism may be triggered by current house price increases. In this section, I extend the model to incorporate the expectation channel.

Suppose a shock in period t changes the house price relative to the steady state, i.e., $\frac{p_t - p_{t-1}}{p_{t-1}} \equiv g_p > 0$. In Section 7.1, households expect future house prices to stay at the current level, i.e., $\mathbb{E}_t p_{t+k} = p_t$, $k = 1, 2, \ldots$. To model the expectation channel, I assume that expected future house prices evolve according to

$$\mathbb{E}_t p_{t+k} = (1 + g_p \nu^k)\mathbb{E}_t p_{t+k-1}, \quad k = 1, 2, ..., \tag{15}$$

where ν is the parameter governing the extent to which households form their expectations based on current house price growth. $\nu = 0$ corresponds to a "flat" expected path of house prices as in Section 7.1, whereas $\nu = 1$ corresponds to the case of exponential growth.

Figure G1 shows the responses for different values of ν. Compared to $\nu = 0$ (Section 7.1), allowing a rising path of expected house prices increases the borrowing propensity and housing investment. The average consumption response, in contrast, stays roughly the same, because the cut back of young households is compensated by higher consumption of old households. The analysis confirms that expected growth in house prices stimulates housing investment. As long as ν is relatively small, the responses are similar to the baseline expectation scenario.

Figure G1: Responses to 1% increase in house price

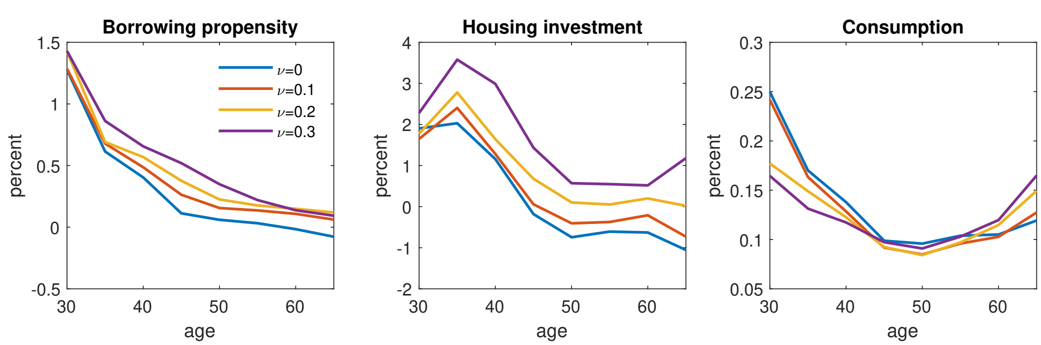

H Additional Figures

Figure H1: Time series fed into the model

Notes: The real house price is constructed by deflating the FHFA house price index by the CPI. The real mortgage rate is constructed by subtracting the inflation from the average 30-year fixed mortgage rate. The credit condition from 2000 to 2007 is measured by the change in the single-family conforming loan limit.