Journal of Monetary Economics 60 (2013) 221–238

Contents lists available at SciVerse ScienceDirect

Journal of Monetary Economics

journal homepage: www.elsevier.com/locate/jme

Housing and debt over the life cycle and over the business cycle

Matteo Iacoviello [a,*], Marina Pavan [b]

[a] Division of International Finance, Federal Reserve Board, 20th and C St. NW, Washington, DC 20551, United States
[b] Department of Economics and LEE, Universitat Jaume I, Spain

ARTICLE INFO

Article history:
Received 27 September 2010
Received in revised form
23 October 2012
Accepted 23 October 2012
Available online 14 November 2012

ABSTRACT

Housing and mortgage debt are studied in a quantitative general equilibrium model. The model matches wealth distribution, age profiles of homeownership and debt, and frequency of housing adjustment. Over the cycle, the model matches the cyclicality and volatility of housing investment, and the procyclicality of debt. Higher individual income risk and lower downpayments can explain the reduced volatility of housing investment, the reduced procyclicality of debt, and part of the reduced volatility of GDP. In an experiment that mimics the Great Recession, countercyclical financial conditions can account for large drops in housing activity and debt following large negative shocks.

Published by Elsevier B.V.

1. Introduction

What are the business cycle and the life-cycle properties of housing and debt? To answer this question, this paper introduces housing in an equilibrium business cycle model where a house is a lumpy item, can be owned or rented, and can be used as collateral for loans. At the cross-sectional level, the model reproduces the wealth distribution, and replicates the life-cycle profiles of housing and nonhousing wealth. The young, the old and the poor are renters and hold few assets; the middle-aged and the wealth-rich are homeowners. For a typical household, the asset portfolio consists of a house and a large mortgage. The model also reproduces the microeconomic evidence on housing adjustment: homeowners change house size infrequently but in large amounts when they do so; renters change house size often, but in smaller amounts. Over the business cycle, the model replicates two empirical characteristics of housing investment: its procyclicality and its high volatility. In addition, the model matches the procyclicality of mortgage debt. To our knowledge, no previous model with rigorous micro-foundations for housing demand has reproduced these regularities in general equilibrium.

We use the model to look at the role of the housing market in two events of the recent U.S. macroeconomic history: the Great Moderation and the Great Recession. Our choice is motivated by two sets of observations: the reduction in aggregate volatility of the early 1980s occurred when idiosyncratic volatility began to rise and when downpayment requirements were relaxed.[1] The sharp decline in housing investment in the 2007–2009 period occurred when, for a variety of causes, financial conditions became very tight.

Debt and housing in the great moderation. Higher risk and lower downpayments are potentially important determinants of housing demand and housing tenure: higher risk should make individuals reluctant to buy large items that are costly to sell in bad times; lower downpayments should encourage and smooth housing demand. Their role could be relevant given two observations on the post-1980 s period (Fig. 1 and Table 1). First, the volatility of housing investment has fallen more

* Corresponding author. Tel.: +1 202 452 2426.
E-mail address: matteo.iacoviello@frb.gov (M. Iacoviello).

[1] Campbell and Hercowitz (2005) and Gerardi et al. (2010) discuss the role of financial reforms, and Dynan et al. (2007) discuss the evolution of household income volatility.

0304-3932/$ - see front matter Published by Elsevier B.V.
http://dx.doi.org/10.1016/j.jmoneco.2012.10.020

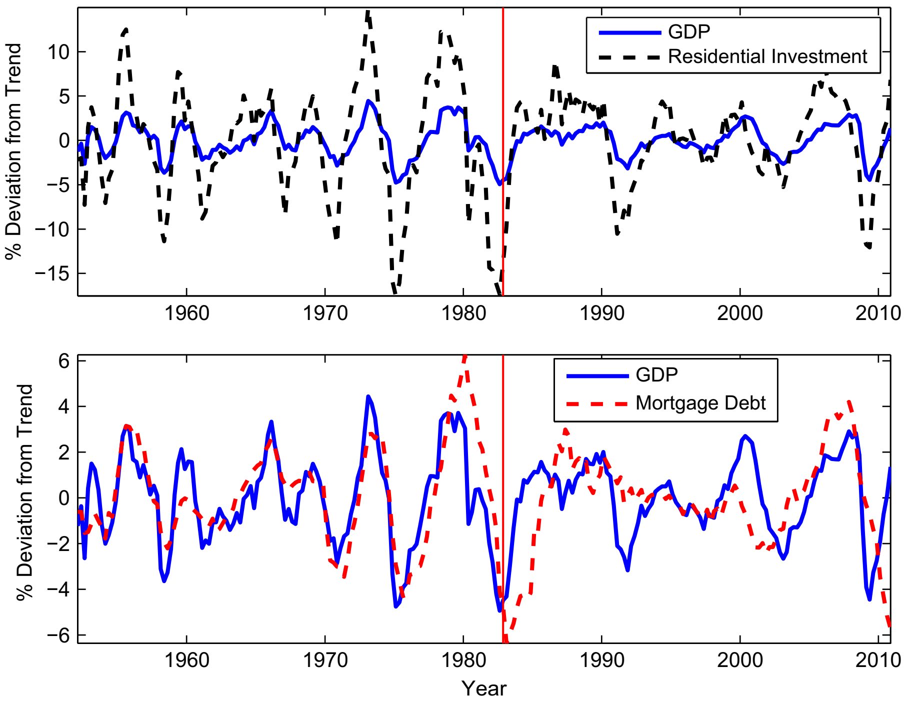

Fig. 1. Mortgage debt, housing investment and GDP. *Note*: Variables are inflation-adjusted, HP-filtered ($\lambda = 1600$).

Table 1
U.S. Economy. Cyclical statistics and housing market facts.

	Early period 1952.I–1982.IV	Late period 1983.I–2010.IV	Whole sample 1952.I–2010.IV
Standard deviation			
GDP	2.09	1.62	1.88
C	0.93	0.83	0.88
IH	7.12	4.45	6.00
IK	4.90	5.36	5.11
Debt	2.23	2.20	2.21
Hours	1.60	1.37	1.49
Housing turnover	0.54 (68.I-82.IV)	0.29	0.40
Correlations			
IH, GDP	0.89	0.75	0.84
Debt, GDP	0.78	0.43	0.63
Hours, GDP	0.82	0.86	0.83
Turnover, GDP	0.69	0.10	0.46
IH, IK	0.36	0.40	0.36
Debt, *C*	0.72	0.37	0.56
Averages			
Homeownership (%)	64	66	65
Debt to GDP (%)	34	59	46
Housing turnover (%)	3.9	4.3	3.2
Gini wealth	0.79	0.83	0.81
Gini labor income	0.40	0.46	0.83
Gini consumption	0.23	0.26	0.25

Notes: *C*, IH and IK are consumption, residential fixed investment and business fixed investment respectively, divided by the GDP deflator (sources: BEA). *GDP* is the sum of the three series. Durables expenditures are included in *IH*. *Debt* is the stock of Home mortgages held by households and nonprofit organizations (source: Flow of Funds Accounts), divided by the GDP deflator. *Hours* are total hours worked for the entire economy from Francis and Ramey (2009). Cyclical statistics (standard deviations and correlations) for all series refer to the series logged and detrended with HP-filter (smoothing parameter 1600). Data on inequality are from Wolff, 2010 (wealth); http://www.census.gov/hhes/www/income/data/ (income); and from Krueger and Perri, 2006 (consumption). Housing turnover is the ratio of total home sales divided by the existing housing stock (see text for the source).

than proportionally relative to GDP; second, the correlations between mortgage debt and GDP and mortgage debt and aggregate consumption have roughly halved, from 0.78 to 0.43 and from 0.72 to 0.37 respectively. In line with the data, lower downpayments and larger idiosyncratic risk reduce the volatility of housing investment, and reduce the correlation between mortgage debt and economic activity. Lower downpayments provide a cushion to smooth housing demand; increase homeownership rates, raising the number of people who do not change their housing consumption over the cycle (relative to an economy with a large number of renters who can become first-time buyers); lead to higher debt, creating a mechanism that weakens the correlation between output and hours. Higher idiosyncratic risk makes wealth–poor individuals more cautious: these individuals adjust consumption, hours, and housing by smaller amounts in response to aggregate shocks. This mechanism is pronounced for housing purchases, since a house is a large item that is costly to purchase and sell; and is reinforced by low downpayments, since low downpayments allow people to borrow more, increasing the utility cost of buying and selling when net worth is lower. Together, lower downpayments and higher risk can explain about 15% of the reduction in the variance of GDP, 60% of the reduction in the variance of housing investment, and the decline in the correlation between debt and economic activity.

Debt and housing in the great recession. During the 2007–2009 period, changes in financial conditions are likely to have made the recession worse. In particular, the housing market appears to have been held back – more than other sectors – by tighter credit conditions and higher borrowing costs. In hindsight, it looks like housing did not stabilize the economy during the recession. We use the model to determine the extent to which housing can smooth regular business cycle shocks but amplify extremely negative ones, by defining "Normal Recessions" as periods of low aggregate productivity, and "Great Recessions" periods of low aggregate productivity coupled with tight credit conditions. An interesting nonlinearity emerges: higher risk and lower downpayments can make housing and debt more stable in response to small positive and negative shocks (as in the Great Moderation), but can make it more fragile in response to large negative shocks (as in the Great Recession).

Previous literature. Two strands of literature study the role of housing in the macroeconomy. On the one hand, business cycle models with housing – Greenwood and Hercowitz (1991), Gomme et al. (2001), Davis and Heathcote (2005), Fisher (2007) and Iacoviello and Neri (2010) – match housing investment well, but abstract from a detailed modeling of the microfoundations of housing demand; these models feature no wealth heterogeneity, no distinction between owning and renting, and unrealistic transaction costs. On the other hand, incomplete markets models with housing – Gervais (2002), Fernandez-Villaverde and Krueger (2004), Chambers et al. (2009), and Díaz and Luengo-Prado (2010) – have a rich treatment of the microfoundations of housing demand, but ignore aggregate shocks: however, because these papers model individual heterogeneity, they are better suited to study issues such as debt, risk, and wealth distribution.

Our model combines both strands of literature. Others have also done so, albeit with a different focus. Silos (2007) studies the link between aggregate shocks and housing choice, but does not model the own/rent decision and assumes convex costs for housing adjustment. Fisher and Gervais (2007) find that the decline in housing investment volatility is driven by a change in the demographics of the population together with an increase in the cross-sectional variance of earnings. Their approach sidesteps general equilibrium considerations. Kiyotaki et al. (2011) study the interaction between borrowing constraints, housing prices, and economic activity. Favilukis et al. (2009) use a two-sector model with housing that also considers the interaction between borrowing constraints and aggregate activity, but address different questions. Finally, Campbell and Hercowitz (2005) study the impact of financial innovation on macroeconomic volatility in a model with two household types. In their model, looser collateral constraints weaken the connection between constrained households' housing investment, debt accumulation and labor supply through a mechanism that shares some features with ours; however, they do not study the interaction between life cycle, risk and housing demand, which are important elements of our story.

We want to be up front about two apparent shortcomings of our approach, especially in light of the 2007–2009 recession, which were mostly imposed on us by the computational limits involved in developing a richer model with more state variables. One limitation is that the model does not endogenize house prices. A second limitation is that the model does not consider mortgage default.[2] We return to and discuss these issues in the concluding section.

2. The model

The economy is a version of the stochastic growth model with overlapping generations of heterogeneous households, extended to allow for housing investment, collateralized debt and a housing rental market. Aggregate uncertainty is introduced in the form of a shock to total factor productivity. Individuals live at most T periods and work until age $\tilde{T} < T$. Their labor endowment depends on a deterministic age-specific productivity and a stochastic component. After retirement, people receive a pension. Each period, the probability of surviving from age a to $a+1$ is χ_{a+1}. Each period a generation is born of the same measure of dead agents, so that the total population, which is normalized to 1, is constant. When an agent dies, he is replaced by a descendant who inherits his assets.

At each point in time, agents differ by their age and productivity; moreover, agents are assumed to differ in their degree of impatience, following a large literature (see Guvenen, 2011) that suggests that preference heterogeneity may be an important source of wealth inequality.

Household preferences and endowments. Households receive utility from consumption c, leisure $\bar{l}-l$ (where $\bar{l}$ is the time endowment), and service flows s from housing, which are proportional to the housing stock owned or rented. The momentary

[2] In Appendix C (online) we sketch a simple modification of our model that allows households to default.

utility function is

$$u(c,s,\bar{l}-l) = \log c + j \log(\theta s) + \tau \log(\bar{l}-l). \tag{1}$$

Above, $\theta = 1$ if $s = h > 0$ (the individual owns), while $\theta < 1$ if $h=0$ (the individual rents). The assumption for θ implies that a household experiences a utility gain when transitioning from renting to owning. Homeowners need to hold a minimum size house $\underline{h}$, and rental units may come in smaller sizes than houses, as in Gervais (2002).[3]

Time supplied in the labor market is paid at the wage w_t. The productivity endowment of an agent at age a is given by $\eta_a z$, where η_a is a deterministic age-specific component and z is a shock to the efficiency units of labor, $z \in \tilde{Z} \equiv \{z^1, \ldots, z^n\}$. The shock follows a Markov process with transition matrix $\pi_{z,z'} = \Pr(z_{t+1} = z' | z_t = z)$ and stationary distribution $\Pi(z) = \Pr(z_t = z)$. The total amount of labor efficiency units $\sum_{i=1}^{n} z^i \Pi(z^i)$ and of age-specific productivity values $\sum_{a=1}^{\tilde{T}} \eta_a \Pi_a$ are constant and normalized to one. From $\tilde{T}+1$ onwards labor efficiency is zero ($z=0$) and agents live off their pension P and their accumulated wealth. Pensions are fully financed through the government's revenues from a lump-sum tax Γ paid by workers. Total net income at age a in period t is denoted by y_{at}. Then:

$$y_{at} = w_t \eta_a z_t l_t - \Gamma \quad \text{if } a \le \tilde{T}; \quad y_{at} = P \quad \text{if } a > \tilde{T}. \tag{2}$$

Households start their life with endowments b_0 and h_0, the accidental bequests left by a dead agent. They can trade a one-period bond b which pays a gross interest rate of R_t. Positive amounts of this bond denote a debt position. Households cannot borrow more than a fraction $m_H < 1$ of their housing stock and a fraction m_Y of their expected earnings:

$$b_t \le \min\{m_H h_t, m_Y \Re_t(y_{at}; R_t, w_t)\}. \tag{3}$$

Above,

$$\Re_t(\bar{y}_{at}; R_t, w_t) = \bar{y}_{at} + \sum_{s=a+1}^{T} \frac{E_t(y_s | \bar{y}_{at}; w_t)}{(R_t)^{s-a}}$$

approximates the present discounted value of lifetime labor earnings and pension.[4] This borrowing constraint prevents the elderly from borrowing too much late in life (when the present discounted value of earnings is low), as in the data. The constraint is also consistent with lending criteria in the mortgage market that take into account minimum downpayments, ratios of debt payments to income, current and expected future employment conditions.[5] Finally, an owner incurs a transaction cost whenever he adjusts the housing stock: $\Psi(h_t, h_{t-1}) = \psi h_{t-1}$ if $|h_t - h_{t-1}| > 0$. This assumption captures common practices in the housing market that require, for instance, fees paid to realtors to be equal to a fraction of the value of the house being sold. Summing up, households maximize expected lifetime utility:

$$E_1 \left(\sum_{a=1}^{T} \beta_i^{a-1} \lambda_a \left(\prod_{\tau=1}^{a-1} \chi_{\tau+1} \right) u(c_a, s_a, \bar{l}-l_a) \right), \tag{4}$$

where E_1 denotes expectations at age $a=1$, λ_a is a deterministic preference shifter that mimics changes in household size, and β_i is a household-specific discount factor. In the calibration, households are born either impatient (low β) or patient (high β).

Financial sector and housing rental market. A competitive financial sector collects deposits from households who save, lends to firms and households who borrow, and buys capital to be rented in the same period to tenants. The financial sector can convert the final good into housing and capital at no cost. This assumption ensures that the consumption prices of housing and capital are constant. Let p_t be the price of one unit of rental services. Then a no-arbitrage condition holds such that the net revenue from lending one unit of financial capital must equal the net revenue from renting one unit of housing capital

$$p_t = 1 - E_t((1-\delta_H)/R_{t+1}) \tag{5}$$

at any t, where δ_H is the depreciation rate of the housing stock.

Production. The goods market is competitive and characterized by constant returns to scale, so one can consider a single representative firm. Output is produced according to

$$Y_t = A K_{t-1}^{\alpha} L_t^{1-\alpha}, \tag{6}$$

where K and L are total capital and labor input; α is the capital share, and $A \in \tilde{A} \equiv \{A^1, \ldots, A^{n_A}\}$ is a shock to total factor productivity. This shock follows a Markov process with transition matrix $\pi_{A,A'} = \Pr(A_{t+1} = A' | A_t = A)$. The aggregate

[3] The log specification over consumption and housing services follows Davis and Ortalo-Magne (2011) who find that, over time and across cities, the expenditure share on housing is constant.

[4] To compute $\Re_t$, we fix interest and wages at current values. To compute $\bar{y}_{at}$, l_t is assumed to equal $\bar{l}$ for $t \le \tilde{T}$.

[5] Our model imposes the borrowing constraint upon purchase and in every period after. Since we do not endogenize house prices, this setup mimics a world where refinancing is costless, except late in life when expected income is low. In real life, changes in house prices can imply an ex post violation of the borrowing constraint given that debt repayments are fixed at the beginning of the loan and are not tied to changes in house prices. Modeling this mechanism would be interesting, but is beyond the scope of this paper.

feasibility constraint requires that production of the good Y_t equals the sum of aggregate consumption C_t, investment in the stock of aggregate capital K_t, investment in the stock of aggregate housing $H_t = H_t^o + H_t^r$, and total transaction costs incurred by homeowners for changing housing stock, denoted by Ω_t:

$$C_t + H_t - (1-\delta_H)H_{t-1} + \Omega_t + K_t - (1-\delta_K)K_{t-1} = Y_t, \tag{7}$$

with δ_H and δ_K denoting the depreciation rates of housing and capital, respectively.

The Household Problem and Equilibrium. Denote with $\Phi_t \equiv \Phi_t(z_t, b_{t-1}, h_{t-1}; \beta, a)$ the distribution of households over earnings shocks, asset holdings, housing wealth, discount factors and ages in period t. Given aggregate volatility, Φ_t changes over time. Following the strategy of Krusell and Smith (1998), agents use one moment of the distribution Φ – the aggregate capital stock K—in order to forecast future prices. As documented in the online Appendix A, using one moment only allows us to obtain a fairly precise forecast, as measured by the R^2 of the forecasting equations, which are between 0.99 and 1.[6]

It is convenient to write the household optimization problem recursively. The individual states are productivity z_t, debt b_{t-1}, and housing wealth h_{t-1}. Agents observe beginning of period capital K_{t-1} and approximate the evolution of aggregate capital and labor with linear functions that depend on the aggregate shock A_t. Denote $x_t \equiv (z_t, b_{t-1}, h_{t-1}, A_t, K_{t-1})$ the vector of individual and aggregate states. The dynamic problem of an age a household is

$$V_a(x_t;\beta_i) = \max_{I^h \in \{0,1\}} \{I^h V_a^h(x_t;\beta_i) + (1-I^h)V_a^r(x_t;\beta_i)\}, \tag{8}$$

where $I^h = 1$ corresponds to the decision to own and $I^h = 0$ to rent. V_a^h and V_a^r are the value functions if the agent owns and rents, respectively, given by

$$V_a^d(x_t;\beta_i) = \max_{c_t, b_t, s_t, l_t} \{\lambda_a u(c_t, s_t, \bar{l} - l_t) + \beta_i \chi_{a+1} \sum_{z', A'} \pi_{A,A'} \pi_{z,z'} V_{a+1}(x_{t+1};\beta_i)\} \tag{9}$$

$$\text{s.t. } c_t + I^h h_t + (1-I^h)p_t s_t + \Psi(h_t, h_{t-1}) = y_{at} + b_t - R_t b_{t-1} + (1-\delta_H)h_{t-1}, \tag{10}$$

$$b_t \le \min\{m_H h_t, m_Y \Re_t\},\ c_t \ge 0,\ l_t \in (0, \bar{l}), \tag{11}$$

$$s_t = h_t > 0 \quad \text{if } I^h = 1, \quad h_t = 0 \quad \text{if } I^h = 0, \quad K_t = F^K(K_{t-1}, A_t), \quad L_t = F^L(K_{t-1}, A_t). \tag{12}$$

for $d = h, r$. Here F^K and F^L are linear functions in K_{t-1}, whose parameters depend on A_t. They denote the law of motion of the aggregate state, which agents take as given.

At the agent's last age, $V_{T+1}(x_{T+1};\beta) = 0$ for any $(x_{T+1};\beta)$. The recursive competitive equilibrium for our economy can then be defined.

Definition 1. An equilibrium consists of value functions $V_a(x_t;\beta)$; policy functions $I_a^h(x_t;\beta)$, $h_a(x_t;\beta)$, $s_a(x_t;\beta)$, $b_a(x_t;\beta)$, $c_a(x_t;\beta)$, $l_a(x_t;\beta)$ for each β, age a and period t, prices R_t, w_t and p_t, quantities K_t, L_t, H_t^o and H_t^r for each t, taxes Γ and pensions P, and laws of motion F^K and F^L such that at any t:

Agents optimize: Given R_t, w_t, p_t, and the laws of motion F^K and F^L, the value functions solve the individual's problem (8), with the corresponding policy functions.

Rental prices are as in (5) and factor prices satisfy:

$$R_t - 1 + \delta_K = \alpha A_t (K_{t-1}/L_t)^{\alpha-1}, \quad w_t = (1-\alpha)A_t(K_{t-1}/L_t)^{\alpha}. \tag{13}$$

In the labor market aggregate labor is equal to $L_t = \int l_a(x_t;\beta)\eta_a z_t \, d\Phi_t$, and the goods market clears as in (7) where H_t and Ω_t are defined as

$$H_t = H_t^o + H_t^r = \int I_a^h(x_t;\beta)h_a(x_t;\beta)\, d\Phi_t + \int (1 - I_a^h(x_t;\beta))s_a(x_t;\beta)\, d\Phi_t, \tag{14}$$

$$\Omega_t = \int \Psi(h_a(x_t;\beta), h_{t-1})\, d\Phi_t. \tag{15}$$

The government budget is balanced, so that $\sum_{a=1}^{\tilde{T}} \Pi_a \Gamma = \sum_{a=\tilde{T}+1}^{T} \Pi_a P$.

The laws of motion for the aggregate capital and aggregate labor are given by (12).

Appendix A provides the details on our computational strategy.

[6] Letting agents use both the aggregate capital stock K and the housing stock H in forecasting future prices produced nearly identical results. Higher moments of the wealth distribution could be relevant in predicting future prices and yield different aggregate dynamics: in that case, our decision rules would describe a bounded rationality equilibrium, rather than a good approximation to the rational expectations equilibrium. Yet the evidence that adding H to the state variables does not change aggregate dynamics leads us to be skeptical of this interpretation. See Young (2010) for a discussion of these issues.

3. Calibration

Our calibration is in Table 2. One period is a year. Agents enter the model at age 21, retire at age 65, and die no later than age 90. The survival probabilities are for men aged 21–90 from the Decennial Life Tables for 1989–1991. Each period, the measure of those who are born is equal to the measure of those who die. The age polynomial λ_a, which captures the effect of demographic variables in the utility, is taken from Cagetti (2003) and approximated using a fourth-order polynomial. The household size is normalized to 1 at age 21, peaks at 2.5 at age 40, and declines thereafter.

The deterministic profile of efficiency units of labor for males aged 21–65 is taken from Hansen (1993) and approximated using a quadratic polynomial: it is normalized at 1 at age 21, peaks at 1.82 at the age of 50, and reaches 1.61 at the age of 65. Upon retirement, an agent receives a pension equal to 40% of the average labor income. The idiosyncratic shock to labor productivity is specified as

$$\log z_t = \rho_Z \log z_{t-1} + \sigma_Z (1-\rho_Z^2)^{1/2} \varepsilon_t, \quad \varepsilon_t \sim Normal(0,1), \tag{16}$$

which is approximated with a three-state Markov process following Tauchen (1986). There is a vast literature on the nature and specification of a parsimonious yet empirically plausible income process: the bulk of the studies (see Guvenen, 2011) look at earnings (rather than wages) and estimate persistence coefficients ranging from 0.7 to 0.95. One exception is Floden and Lindé (2001), who use PSID data to estimate an AR(1) process for wages similar to ours and find an autocorrelation coefficient of 0.91. Based on this evidence, ρ_Z is set at 0.9 (Section 8 contains robustness analysis, based on evidence from other studies). The standard deviation of the labor productivity process is set at $\sigma_Z = 0.30$ (see online Appendix B). Later, σ_Z is increased to 0.45 to capture the increased earnings volatility of the 1990s, and to study the consequences for macroeconomic aggregates of increased risk at the household level, as emphasized by Moffitt and Gottschalk (2008) and Dynan et al. (2007).

There are two classes of households, a "patient" group with a discount factor of 0.999 (one third of the population) and an "impatient" group with a discount factor of 0.941 (two thirds of the population). The high discount factor pins the average real interest rate down to 3%. The low discount factor is in the range of estimates in the literature (see, for instance, Hendricks, 2007). The gap between discount rates and the relative population shares deliver a Gini coefficient for wealth around 0.75, close to the data. Section 8 discusses the properties of the model when all people have identical discount rates. When $\tau = 1.65$ and the endowment of time $\bar{l} = 2.65$, time spent working is 40% of the agents' time.

The weight on housing in utility is set at j=0.15, and the depreciation rate for housing is $\delta_H = 0.05$. These parameters yield average housing investment to private output ratios around 7%, and a ratio of the housing stock to output 1.4. These

Table 2
Parameter values for the benchmark model economy.

	Parameter	Value	Target/source
Preferences			
Discount factor, patients	β_H	0.999	R=3%
Discount factor, impatients	β_L	0.941	Hendricks (2007)
Fraction of impatient agents	–	2/3	Gini coefficient of Wealth: 0.73
Weight on leisure in utility	τ	1.65	–
Productive time	$\bar{l}$	2.65	Time worked: 40%
Weight on housing in utility	j	0.15	H/Y=1.4
Utility, renting vs. owning	θ	0.838	Homeownership rate = 64%
Utility weights (family size)	λ_a	See text	Cagetti (2003)
Life, retirement			
Survival probabilities	Π_a	See text	Decennial life tables
Retirement period	$\tilde{T}$	46	Retirement age 65 years
Pension	P	0.4 × inc.	40% average income
Technology			
Capital share	α	0.26	K/Y=2.2
Capital depreciation rate	δ_K	0.09	IK/Y=0.20
Housing depreciation rate	δ_H	0.05	IH/Y=0.07
Autocorrelation, technology shock	ρ_A	0.925	King and Rebelo (1999)
Standard deviation, technology shock	σ_A	0.0148	$\sigma(Y)$ = 2.09%
Housing transaction cost	ψ	0.05	National Association of Realtors (2005)
Minimum house size	$\underline{h}$	1.5 × inc.	See text
Borrowing			
Max debt, fraction lifetime wage	m_Y	0.25	See text
Maximum debt, fraction of house	m_H	0.75	See text
Individual income process			
Autocorrelation, earnings shock	ρ_Z	0.90	Floden and Lindé (2001)
Standard deviation, earnings shock	σ_Z	0.30	See Appendix B
Age-dependent earnings ability	η_a	See text	Hansen (1993)

values are in accordance with the National Income and Product Accounts and the Fixed Assets Tables.[7] Finally, the housing transaction cost is set at $\psi = 5\%$ based on estimates from the National Association of Realtors (2005). Section 8 conducts robustness analysis for alternative values of ψ and δ_H.

The values of $\alpha = 0.26$ and $\delta_K = 0.09$ yield an average capital to output ratios around 2.2 and average business investment to output ratios around 20%. The aggregate shock is calibrated to match the standard deviation of output in the data for the period 1952–1982, using a Markov-chain specification with seven states that matches the following autoregression for the log of total factor productivity:

$$\log A_t = \rho_A \log A_{t-1} + \sigma_A (1-\rho_A^2)^{1/2} \varepsilon_t, \quad \varepsilon_t \sim Normal(0,1), \tag{17}$$

where $\rho_A = 0.925$ and $\sigma_A = 0.0148$. After rounding, the first number mimics a quarterly autocorrelation rate of productivity of 0.979, as in King and Rebelo (1999). The second number is chosen to match the standard deviation of output to its data counterpart.

The baseline calibration sets the maximum loan-to-value ratio m_H at 0.75 ($m_H = 0.85$ in the late period calibration). The value of m_Y is set at 0.25 in the baseline and raised to 0.5 in the late period: with these numbers, the income constraint only binds late in life, preventing old homeowners from borrowing. Aside from this, the choice for m_Y is of small importance for the model dynamics. Lastly, the minimum-size house available for purchase ($\underline{h}$) costs 1.5 times the average annual pre-tax household income.[8] Together with the minimum house size, the parameter that has a large impact on homeownership is the utility penalty for renting (θ). A value of $\theta = 0.838$ to obtain a homeownership rate of 64%, as in the data for the period 1952–1982.

4. Steady-state results

In this section, we illustrate the steady-state properties of the model.

Housing choices. For renters, there is a threshold amount of liquid assets such that, if assets exceed the threshold, renters become homeowners. Also, the larger initial liquid assets are, the less likely a household is to borrow to finance its housing purchase.

Fig. 2 plots homeowners' choices as a function of initial house and liquid assets. Homeowners can stay put, increase size, downsize or switch to renting. Larger liquid assets trigger larger housing choices, and transaction costs create a region of inaction where the housing stock is constant. If liquid wealth falls, the household either downsizes or switches to renting. One feature of the model is that, for a household with very small liquid assets, the housing tenure decision is non-monotonic in the initial level of housing wealth. Consider, for instance, a homeowner with liquid assets equal to about one. If the initial house size is small, the homeowner does not change house size, since, given the small amount of assets, the house size is closer to its optimal choice. If the initial house is medium-sized, the homeowner pays the adjustment cost and, because of his low liquid assets, switches to renting. If the initial house size is large, it is optimal to downsize, and to buy a smaller house.

Life-cycle profiles. Fig. 3 plots a typical life-cycle profile in our model. We choose an agent with a low discount factor since the behavior of an agent with low assets and often close to the borrowing constraint best illustrates the main workings of the model. The agent starts life as a renter, with little assets and low income. At the age of 22, he receives a positive income shock, saves to afford the downpayment and buys a house a year after. Prior to buying a house, the individual works more: the positive income shock raises the incentive to work; and such incentive is reinforced by need to set resources aside for the downpayment. Following a series of above average income shocks beginning at the age of 32, the agent buys a larger house at the age of 39. This time, in order to afford the larger house, the individual is much closer to his borrowing limit. In particular, while he owns and is close to the borrowing limit, hours move in the opposite direction to wage shocks, rising in bad times (age 42), falling in good times (age 45): such mechanism is explained in detail in the next Section. As retirement approaches, the agent pays back part of the mortgage, and works more. After retirement, at the age of 70, he moves to a small rental unit.

One dimension where it is illustrative to compare the model with the data is the frequency of housing adjustment for homeowners.[9] Using the 1993 Survey of Income and Program Participation, Hansen (1998) reports that the median homeowner stays in the same house for 8 years. Anily et al. (1999) estimate that the average homeowner lives in the same house for 13 years. The corresponding number for our model is 15 years.

Fig. 4 compares the age profiles of housing, debt, homeownership and hours worked with their empirical counterparts. Overall, the model is able to capture the profiles of these variables, with some wrinkles: for mortgage debt, the model

[7] The Fixed Asset Tables indicate depreciation rates for housing ranging from 1.2 to 4.5%, depending on the type of structure and its use (see Fraumeni, 1997). A slightly higher value accounts for unmeasured labor time that is used to repair, renovate, or maintain or improve the quality of housing at a given location; because higher values are typically considered in the existing literature, especially when housing is broadly interpreted to include consumer durables (Chambers et al., 2009; Gervais, 2002; Díaz and Luengo-Prado, 2010); and because a higher depreciation rate (5% instead of 2%, say) reduces the extent to which aggregate housing tends to decrease on impact following a positive aggregate technology shock in a model with two capital goods.

[8] In the 2009 American Housing Survey only 20% of owner-occupied units have a ratio to current income less than 1.5.

[9] In the model, renters change their housing position every period, since they face no cost in doing so. This assumption is in line with the data, that show that on average renters move about every two years.

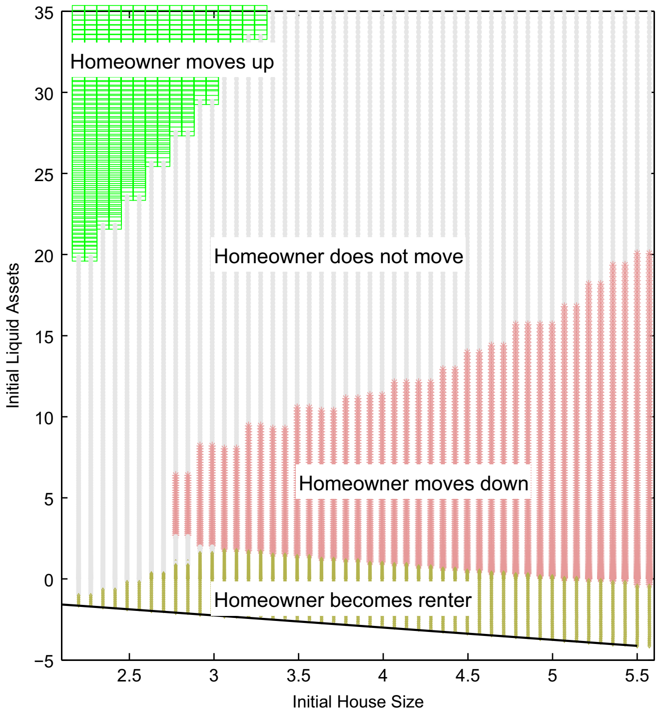

Fig. 2. Homeowner's housing choice as a function of initial house size and liquid assets. *Note:* The figure illustrates the homeowner's housing policy as a function of initial house size and liquid assets. It is plotted for a patient agent who is 65 years old, when aggregate productivity and the average capital labor ratio are equal to their average value. The downward sloping line plots the borrowing constraint.

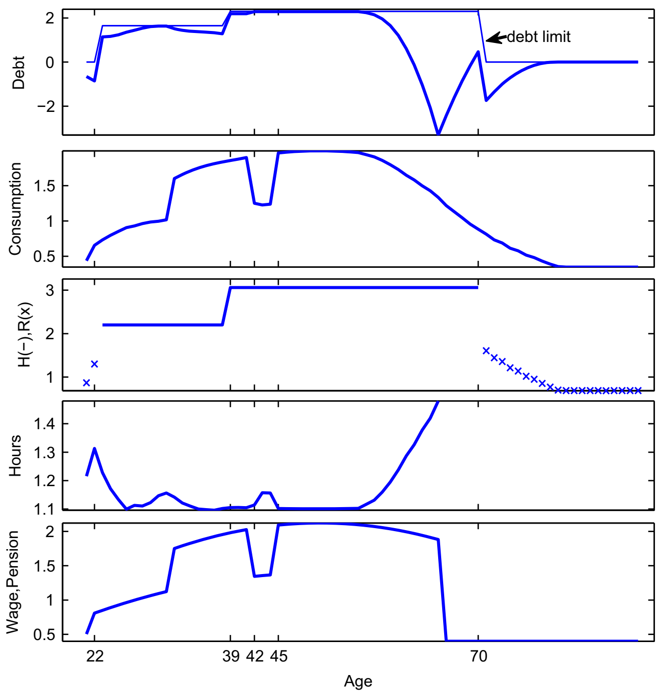

Fig. 3. A Typical life-cycle profile. *Note*: This figure plots life-cycle choices of a randomly chosen impatient agent from birth (age 21) to death (age 90). In panel 1, the thin line denotes the maximum debt limit given the housing choice. In panel 3, the "x" symbol denotes the amount rented when the individual is renting, whereas the solid line denotes the amount owned when the individual owns a house.

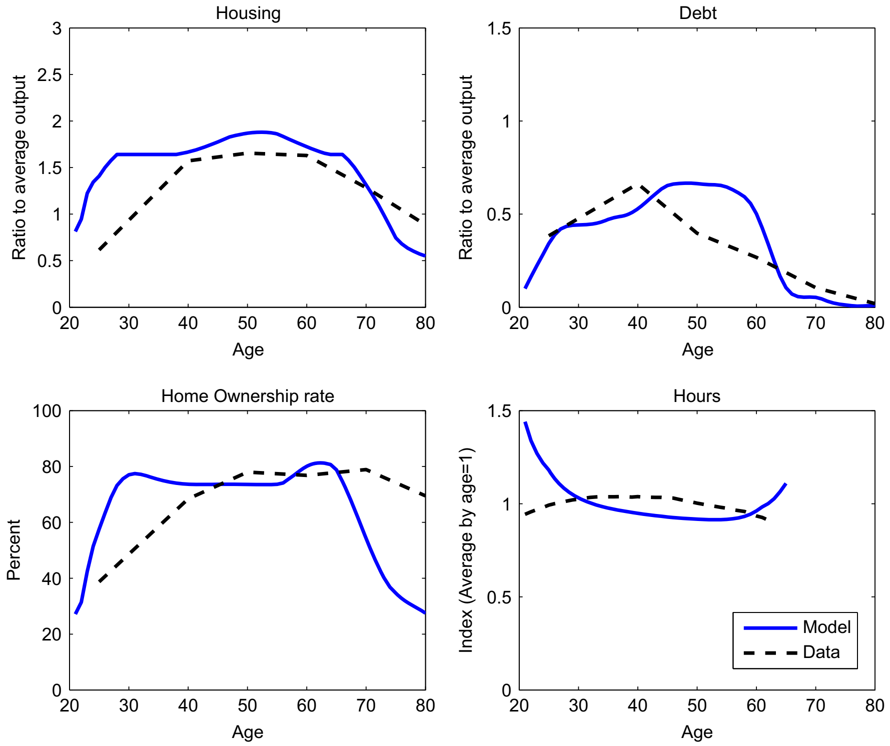

Fig. 4. Comparison between model and data. *Note*: Data for housing, debt and homeownership are from the 1983 Survey of Consumer Finances. For each age, the model variable is the product of the fraction of households in that age holding housing or debt, times the median holding of housing or debt. The data variable is constructed in the same way. Data on hours are for head-of-household males in the PSID (years 1968–1996).

slightly underpredicts debt early in life, and overpredicts debt later in life. The model also underpredicts homeownership later in life: late in life, the absence of any bequest motive and the need to finance consumption expenditure by selling the house more than offset the adjustment costs, thus generating a decline in homeownership. As for hours, our model predicts that hours fall throughout the life cycle, before increasing slightly prior to retirement, whereas the pattern in the data shows a hump-shaped profile of hours with a peak around the age of 35. The model behavior is mostly driven by the borrowing constraint, which prevents agents from smoothing expenditure for low levels of wages, and forces them to reduce consumption and leisure when income is low (which happens at the beginning and the end of the working life), thus raising hours worked when young and when old.

Another dimension where it is instructive to compare the model with the data is the behavior of hours worked by housing and debt status. In the model, young owners with high mortgages work more than owners with low mortgages: between the age of 25 and 45, for instance, borrowers with a mortgage in the top third of the debt distribution work 6% more than those with a mortgage in the bottom third. This prediction of the model finds some indirect support in the studies that have looked at how labor force participation (mostly of married women) varies with housing tenure and mortgage status. For instance, Bottazzi et al. (2007) find, using British data, that women in households that own their homes, and particularly those with greater mortgages, supply more labor. Using Canadian data, Fortin (1995) finds that the participation rate of women in home-owning families with mortgages is 10 percentage points higher than among women from families who either rent or own a home with no mortgage. We return to this issue in Section 5.

The wealth distribution and liquidity constraints. Our model reproduces the U.S. wealth distribution well. Wealth in the model is defined as the sum of housing and financial wealth, that is $\Psi \equiv h - b$. The Gini coefficient in the model is 0.73, and is about the same as in the data (equal to 0.79). The model still underpredicts wealth inequality at the very top of the distribution, both for housing and for total wealth. However, the model does well at matching the fraction of wealth (both housing wealth and overall wealth) held by the poorest 40% of the U.S. population, which has essentially no assets and no debt. Instead, a model without preference heterogeneity would do much worse: as shown in Section 8, the Gini coefficient for wealth in the model with a single discount factor is 0.53, much lower than in the data.

Besides reproducing the overall wealth distribution, the model predicts a mortgage debt to GDP ratio that is in line with the data and a fraction of liquidity-constrained agents that is consistent with available estimates. We use a definition of liquidity constraints based on how close individuals are to their debt limit, and compare it to its data counterpart. Net liquid assets are defined in the model as $\Upsilon \equiv m_H h - b$: using this definition, liquid assets are low whenever agents are close to hitting the borrowing limit. A model agent is taken to be liquidity-constrained if the holdings of Υ are less than two months (16.67% on a yearly basis) of income, following Hall (2011). Using this standard, 45% of households are liquidity constrained in the baseline.[10] We then construct the data analog using the Survey of Consumer Finances (SCF), with the obvious caveat that several assets are in practice quite illiquid and not useful to smooth consumption in response to shocks. Depending on the assets that are considered ''liquid,'' liquidity-constrained households in the SCF range from 32% (using a very broad definition of liquid assets) to 73% (using a narrower definition),[11] with an intermediate value of 51% that obtains when home equity is not considered liquid.[12] Incidentally, it is useful to note that liquidity constraints have been defined differently in different papers: when other indicators have been used, not based on observable measures of wealth but more directly measuring households' access to credit (as in Jappelli, 1990; Gross and Souleles, 2002; Johnson and Li, 2010), smaller evidence of borrowing constraints has been found.

5. Business cycle results

Two aspects of heterogeneity matter for aggregate dynamics: one is exogenous, and reflects the assumption that individuals have different abilities, planning horizons, and utility weights. Other papers have studied these features in life-cycle models with aggregate shocks, so we do not explore them in detail here. The focus here is on the endogenous component of heterogeneity, which reflects that individuals with different income histories accumulate different amounts of wealth over time; in turn, heterogeneity in wealth implies different individual responses to the same shock.

Workings of the model. Movements in hours are the key propagation mechanism in models with technology shocks as sources of aggregate fluctuations. Consider a stripped-down version of the budget constraint of a working individual that keeps wealth constant between two periods: $b_t = b_{t-1}$ and $h_t = h_{t-1}$.[13] Abstracting from taxes and pensions, this implies the following budget constraint:

$$c_t = w_t \eta_a z_t l_t + \xi_t, \tag{18}$$

where $\xi_t = -(R_t - 1)b_{t-1} - \delta_H h_{t-1}$ measures the resources besides wages that can be used to finance consumption[14]: the term $(1-R)b$ is net interest income; the term $\delta_H h$ is the maintenance cost required to keep housing unchanged. Different values of ξ map into different positions of the agents along the wealth distribution. For a wealthy homeowner (negative b), ξ is positive and large, and wage income is a small fraction of consumption c. For a renter, $h=0$; in addition, assuming that the renter is not saving, $b=0$, so that $\xi = 0$ too. For a homeowner with a mortgage (positive b), ξ is negative. Normalize $\eta_a = 1$ and set aside idiosyncratic shocks, so that $z_t = 1$ at all times. Assuming that ξ stays constant, the log-linearized budget constraint becomes, denoting with $\widehat{x} \equiv (x_t - x)/x$, where x is the steady-state value of a variable:

$$\widehat{c} = \frac{wl}{c}(\widehat{w} + \widehat{l}). \tag{19}$$

This constraint can be interpreted as an equation dictating how much the household needs to work to finance a given consumption stream, given the wage. The larger the desired consumption $\widehat{c}$, the larger the required hours $\widehat{l}$ needed to finance the consumption stream, with an elasticity of hours to consumption given by consumption–wage income ratio $(c/wl) \equiv \phi$. For a wealthy individual, ϕ is high and larger than one, since labor income is a small share of total earnings; for a renter without assets, $\phi = 1$; for an indebted homeowner, $\phi < 1$, reflecting the need to use part of the earnings to finance maintenance costs and to service the mortgage. In other words, a wealthy person needs to increase hours by more than 1% to finance a 1% rise in consumption, since labor income is less than consumption; an indebted homeowner needs to increase hours by less than 1% to finance a 1% rise in consumption, because of the leverage effect; a renter without assets needs to increase hours 1 for 1 with consumption.

The other key equation determining hours is the standard labor supply schedule. Letting ζ denote the steady-state Frisch labor supply elasticity, this curve reads as

$$\widehat{l} = \zeta(\widehat{w} - \widehat{c}). \tag{20}$$

Combining Eqs. (19) and (20) yields

$$\widehat{l} = \zeta\left(\frac{\phi - 1}{\phi + \zeta}\right)\widehat{w}. \tag{21}$$

[10] The baseline model predicts that 70% of renters and 31% of homeowners are liquidity constrained; and that 67% of impatient agents and 2% of patient agents are liquidity constrained.

[11] Our numbers are averages across the 1989, 2001 and 2007 surveys. The 73% figure defines net liquid assets as the sum of cash, checking and saving accounts and other liquid accounts, net of credit card balances. The 32% figure defines liquid assets as *all* financial assets net of credit card debt (excluding only ''quasi-liquid retirement accounts'' such as IRAs and thrift-type accounts) and adding potential home equity loans, defined as the $\max(0, m_H h - d)$, where $m_H = 0.8$, h is the value of the primary residence, and d is any mortgage debt.

[12] In the model, houses both liquid and illiquid: they are illiquid because they can only be bought and sold by paying a large transaction cost, but are also liquid because they can be readily used for collateral. Whether home equity is liquid or not in practice is an open question.

[13] Optimal decisions involve the joint choice of (1) consumption, (2) housing, (3) debt and (4) hours worked. By assuming that housing and debt remain constant across two subperiods, one can study the joint determination of consumption and hours by focusing on the budget constraint and the Euler equation for labor supply only.

[14] Renters have constant shares of housing and nonhousing consumption, so that $c_t = (w_t \eta_a z_t l_t + \xi_t)/(1+j)$, where j is housing preference parameter. With minor modifications, the arguments in the text carry over to this case.

Take the wage as the exogenous driving force of the model, since an exogenous rise in productivity exerts a direct effect on the wage. Whether the rise in the wage leads to an increase in hours depends on whether the consumption–wage income ratio, ϕ, is smaller or larger than one. In other words, all else equal, borrowers ($\phi < 1$) are more likely to reduce hours following a positive wage shock, whereas savers ($\phi > 1$) are more likely to increase them.[15]

If individual labor schedules were linear in wealth, the aggregate labor supply response would be linear in average wealth too, and the wealth distribution would not affect labor supply. Two main forces, however, undo this linearity. First, retirees do not work, so any wealth transfer to and from them may affect how workers respond to wage shocks. Second, the interaction between borrowing constraints and housing purchases creates an interesting nonlinearity. The discussion above assumes that households do not change wealth composition in response to a shock in the wage. However, if households switch from renting to owning (or if they increase their house size) in good times, they typically need to save for the downpayment. This increases the incentive to work: intuitively, if the individual wants to keep consumption constant when he buys the house, he needs to work more hours. This effect creates comovement between hours and housing purchases.[16] In particular, it reinforces the correlation between hours and housing demand in periods when a large fraction of the population has, all else equal, low net worth.

Business cycle statistics. In HP-filtered U.S. data, the variability of housing investment is large, with a standard deviation that is between three and four times that of GDP (in the period 1952–1982). Also, housing investment is procyclical, with a correlation with GDP of 0.89. Together, these two facts imply that the growth contribution of housing investment to the business cycle is larger than its share of GDP. Household mortgage debt is strongly procyclical from 1952 to 1982, but it becomes less procyclical after, with a correlation with GDP that drops from 0.78 to 0.43. Table 3 compares the benchmark model with the data. Overall, our baseline model does a good job in reproducing the relative volatility of each component of aggregate demand. In particular, it can account for about three quarters of the variance of housing investment. On the contrary, the model overpredicts the volatility of aggregate consumption. The volatility of business investment is only slightly lower than in the data. As in many RBC models without an extensive margin of work and without direct shocks to the labor supply, our model underpredicts the volatility of hours (0.33% in the model, 1.6% in the data).

Turning to debt, the model does well in reproducing its cyclical behavior.[17] The key to this result is that the bulk of the debt holders (mostly impatients) upgrades housing in good times by taking out a (larger) mortgage. At the same time, the model overpredicts the volatility of debt itself: the standard deviation of the model variable is about four times larger than in the data. One reason for the higher volatility of debt in the model has to do with the simplifying assumption that only one financial asset is available, whereas in the data some households (especially the wealthy) own simultaneously a mortgage and other financial assets. If debt of low-wealth households is more volatile than debt of high-wealth households, our model variable can exhibit more volatility than its data counterpart.

One dimension where it is useful to compare the model with the data pertains to home sales. In our model, a sale is every instance in which a household pays the transaction cost to change its housing: this involves own-to-own, rent-to-own and own-to-rent transitions. By this metric, the average turnover rate in the model (the ratio of sales to total houses) is 4%, a number that matches the 3.9% in the data.[18] Moreover, the model correlation between turnover rate and GDP is 0.39, and the standard deviation is 0.29. The corresponding numbers from the data are 0.69 and 0.54. The positive correlation between sales and economic activity that the model captures reflects the presence of liquidity constraints: when the economy is in recession and household balance sheets have deteriorated, the potential movers in the model find their liquidity so impaired, whether they are owners or renters, that they are better off staying in their old house rather than attempting to move and paying the transaction cost.

6. Effects of lower downpayments and higher risk

Having shown above that the model broadly captures business cycles, this section considers two experiments. In the first, the downpayment is lowered from 25% to 15%. In the second, households' income risk increases, and the standard deviation of income σ_Z rises from 0.30 to 0.45. Our experiment is intended to mirror two of the main changes that have occurred in the U.S. economy since the mid-1980s. The model results are in Table 4.

A decline in downpayments. Lower downpayments (column 2 in Table 4) lead to an increase in the homeownership rate (from 64% to 76%) and to a higher level of debt (from 31% to 50% of GDP). Smaller downpayments allow more housing ownership among the portion of the population with very little net worth. While debt is higher, the increase in homeownership works to keep total wealth inequality unchanged: financial wealth inequality is higher, but housing

[15] Ziliak and Kniesner (1999) find that male labor supply elasticities increase with wealth. Yoshikawa and Ohtaka (1989), using Japanese data, find that labor supply is an increasing function of the wage for families that own or plan to own, and a decreasing function of the wage for families that rent.

[16] The limiting case of zero forced savings would be the case in which no downpayment is needed to buy a house. In that case the individual can keep consumption constant at the time of the purchase without increasing hours worked if transaction costs are zero. If the individual has to pay the transaction cost, this provides an incentive to work more at the time of the purchase. Campbell and Hercowitz (2005) propose a similar argument to discuss the relationship between hours and durable purchases.

[17] Household debt is defined $D_t = \int_{b>0} b_a(x_t;\beta)\, d\Phi_t$ (that is, the average of the household liabilities).

[18] The turnover rate in the data is constructed as the sum of sales of existing single-family homes (source: National Association of Realtors) plus new single-family homes sold (from Census Bureau), divided by the total housing stock (from Census Bureau). The series starts in 1968.

Table 3
U.S. economy and baseline model. Comparison for the early period.

	1952.I–1982.IV (early period)	Model
Standard deviation		
GDP	2.09	2.09
C	0.93	1.63
IH	7.12	6.42
IK	4.90	4.16
Debt	2.23	8.34
Hours	1.60	0.33
Housing turnover	0.54 (68.1–82.IV)	0.29
Correlations		
IH, GDP	0.89	0.66
Debt, GDP	0.78	0.71
Hours, GDP	0.82	0.65
Turnover, GDP	0.69	0.39
IH, IK	0.36	0.18
Debt, C	0.72	0.85
Averages		
Homeownership (%)	64	64
Debt to GDP (%)	34	31
Housing turnover (%)	3.9	4.0
Gini wealth	0.79	0.73
Gini labor income	0.40	0.41
Gini consumption	0.23	0.26
Liquidity constrained	0.51	0.45

Notes: The model moments are based on statistics from a simulation of 5000 periods. Liquidity constrained agents are those who own liquid assets less than 16.67% (two months in a year) of annual income.

Table 4
Model predictions, changing downpayment requirements and income volatility.

	(1) Baseline Early period $m_H = 0.75$, $\sigma_Z = 0.3$	(2) $m_H = 0.85$, $\sigma_Z = 0.3$	(3) $m_H = 0.75$, $\sigma_Z = 0.45$	(4) Late period $m_H = 0.85$, $\sigma_Z = 0.45$
Standard deviation				
GDP	2.09	2.08	2.05	2.03
C	1.63	1.63	1.66	1.68
IH	6.42	5.94	5.52	5.04
IK	4.16	4.05	4.21	4.16
Debt	8.34	3.04	2.61	1.44
Hours	0.33	0.32	0.31	0.31
Housing Turnover	0.29	0.44	0.21	0.21
Correlations				
IH, GDP	0.66	0.69	0.55	0.54
Debt, GDP	0.71	0.63	0.50	0.39
Hours, GDP	0.65	0.64	0.47	0.42
Turnover, GDP	0.39	0.77	0.42	0.28
IH, IK	0.18	0.24	0.08	0.09
Debt, C	0.85	0.77	0.68	0.58
Averages				
Homeownership (%)	64	76	59	67
Debt to GDP (%)	31	50	23	35
Housing turnover (%)	4.0	3.0	5.1	5.6
Gini wealth	0.73	0.73	0.73	0.73
Gini labor income	0.41	0.41	0.48	0.48
Gini consumption	0.26	0.26	0.31	0.31
Liquidity constrained	0.45	0.45	0.39	0.38

Notes: Baseline calibration and sensitivity analysis. (1) is the baseline calibration that is targeted to the U.S. data for the period 1952–1982. In (2), the loan-to-value ratio rises from 0.75 to 0.85. In (3), increase earnings volatility rises from 0.3 to 0.45. In (4), both loan-to-value ratio and earnings volatility are increased to match the U.S. economy for the period 1983–2010.

wealth inequality is lower. Turning to business cycles, the rise in m_H tends to reduce the volatility of housing investment, from 6.42% to 5.94%, for two reasons. The first reason has to do with adjustment costs: on average, because of adjustment costs, homeowners modify their housing little over time relative to renters. The second motive operates through the interaction of labor supply and housing purchases. As explained above, indebted homeowners are more likely, compared to renters, to reduce hours in response to positive technology shocks, so their presence dampens aggregate shocks. Therefore, the higher homeownership rate induced by looser borrowing constraints reduces aggregate volatility.[19]

An increase in individual earnings volatility. Column 3 in Table 4 shows that, following a rise in σ_Z, the homeownership rate falls from 64% to 59%: higher risk makes individuals more reluctant to buy an asset that is costly to change. All else equal, the lower homeownership rate would tend to increase the volatility of housing investment, since renters change housing consumption more often. However, this effect is more than offset by the behavior of those who remain homeowners: these people are now more reluctant to change their housing consumption (relative to a world with less individual risk). This occurs because modifying housing, in the presence of transaction costs, depletes holdings of liquid assets and increases the utility cost of a negative idiosyncratic shock, thus increasing the option value of not adjusting the stock for given changes in net worth. Quantitatively, the higher earnings volatility reduces the standard deviation of housing investment from 6.42% to 5.52%. Moreover, higher income volatility also reduces the sensitivity of debt to aggregate shocks, since debt is used to finance housing purchases, and housing purchases respond less to shocks.

Combining lower downpayments and higher volatility. The last column of Table 4 shows the effects of combining lower downpayments and higher volatility. The two forces together predict an increase in homeownership from 64% to 67%. The data counterpart is a rise from 64% to 66%. Moreover, the joint effect of these two forces makes debt less procyclical, as in the data. The correlation between debt and output falls from 0.71 to 0.39, a change that is remarkably similar to the data (from 0.78 to 0.43, see Table 1).[20] Together, lower downpayments and high idiosyncratic volatility reduce the standard deviation of GDP from 2.09% to 2.03%, and the standard deviation of housing investment from 6.42% to 5.04%. When these numbers are compared to the data, the two changes combined can account for 13% of the variance reduction in GDP and about 60% of the variance reduction in housing investment. The model's explanation for this result is as follows: with lower downpayments and higher income volatility, leveraged households become more *cautious* in response to aggregate shocks, thus changing their debt and housing position by less when aggregate productivity changes. This is especially true for housing, relative to other categories of expenditure, since housing is a highly durable good and is subject to adjustment costs. Because individuals are reluctant to adjust their housing consumption during uncertain times, the sensitivity of hours to aggregate shocks falls too. As a consequence, total output is less volatile.

Fig. 5 offers an illustration, plotting the model dynamics when technology switches from its average value to a higher value (a 1% rise) in period 1. The responses are larger in the earlier period. On impact, housing falls before rising strongly in period 1. This result is well known in the household production literature (see, for instance, Fisher, 2007). In models with housing and business capital, business capital is useful for producing more types of goods than housing capital. Hence, after a positive productivity shock, the rise in the marginal product of capital implies that there is a strong incentive to move resources out of the housing to build up business capital, and only later is housing accumulated. The key aspect to note here is that higher idiosyncratic risk and lower downpayment requirements dampen the incentive to adjust housing capital, so that housing investment becomes less volatile.

Our result that higher individual uncertainty reduces the volatility of aggregate housing investment echoes the results of papers that study how durable purchases respond to changes in income uncertainty in (*S*,*s*) models resulting from transaction costs. Eberly (1994), using data from the Survey of Consumer Finances, considers automobile purchases in presence of transaction costs: she finds that higher income variability broadens the range of inaction, and that the effect is larger for households that are liquidity constrained. Foote et al. (2000) find a similar result using data on car holdings from the Consumer Expenditure Survey, and offer an explanation that involves the presence of liquidity constraints and precautionary saving: adjusting the capital stock for people with low levels of net worth depletes holdings of liquid assets and increases the utility cost of a negative idiosyncratic shock, thus increasing the option value of not adjusting the stock for given changes in net worth.

7. Debt and housing in a great recession experiment

The finding that housing and debt are less sensitive to aggregate shocks when downpayments are low and idiosyncratic risk is high can account for part of the Great Moderation, but is at odds with the events of the 2007–2009 financial crisis, when both housing and debt fell substantially. Explaining the crisis is beyond the scope of this paper, but this section shows that our model – expanded to take into account the "credit crunch" – can generate, at least qualitatively, the observed response of housing and debt in the Great Recession. We extend the stochastic structure of the model so that, when the worst technology shocks hit, credit standards get tighter too, in the form of lower loan-to-value ratios and higher

[19] Campbell and Hercowitz (2005) show that financial innovations explain more than half of the reduction in aggregate volatility in a model with downpayment constraints. Aside from modeling differences (our model considers the owning/renting margin and addresses issues related to life cycle, lumpiness and risk that are absent in their setup), the intuition they offer for their result carries over to our model, but the effect of lower downpayment requirements is quantitatively smaller. The differences might depend on one modeling assumption: in our setup, indebted homeowners mitigate aggregate volatility, but this effect is partly offset by the wealthier homeowners (the creditors) who tend to increase aggregate volatility by working relatively more in response to positive aggregate shocks; instead, Campbell and Hercowitz assume that labor supply of wealthy homeowners is constant, thus killing this offsetting mechanism.

[20] The correlation between debt and consumption falls in the model from 0.85 to 0.58, a decline similar to the data (from 0.72 to 0.37).

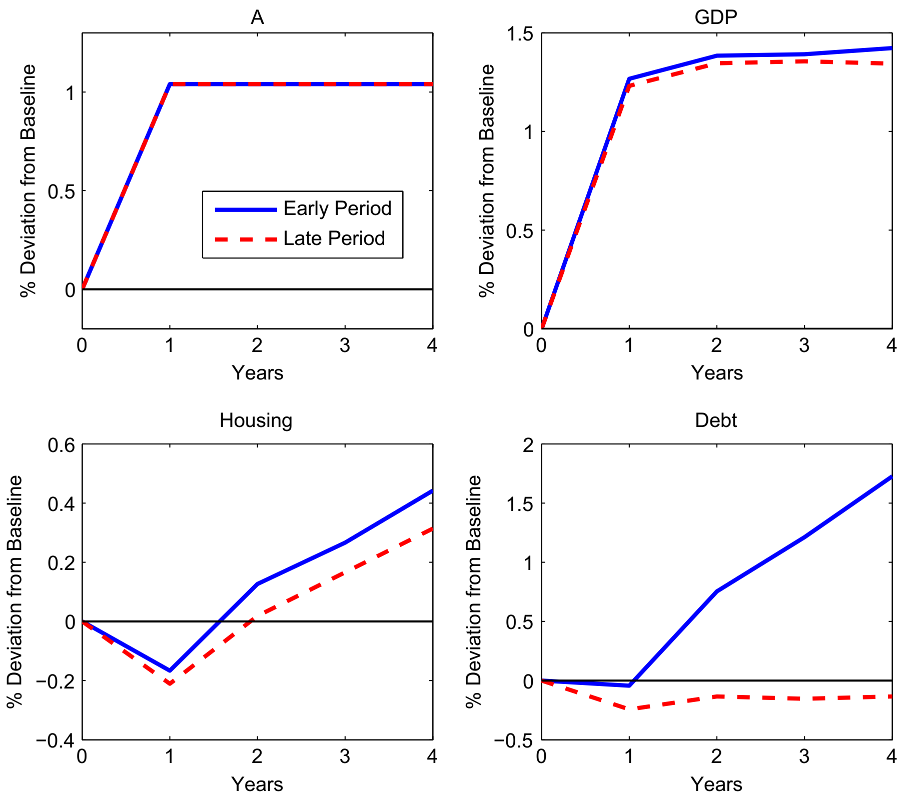

Fig. 5. Responses to a positive technology shock: early and late period calibration. *Note*: Model dynamics following an exogenous switch in aggregate productivity A (in period zero) from the median state to next higher value (a 1% increase) lasting four periods.

costs of financial intermediation (higher borrowing interest rates). In other words, consistent with the post-2007 evidence,[21] recessions are now a combination of negative financial and negative technology shocks occurring simultaneously. This scenario is implemented by assuming that the maximum loan-to-value ratio m_H changes over time as a function of total factor productivity, A_t: formally, $m_{H,t} = m_H(A_t)$. Moreover, an additional cost of financial intermediation is introduced in the form of an interest rate premium $r_t^p = r^p(A_t)$ to be paid by debtors. The budget constraint for a home buyer become respectively:

$$c_t + h_t + \Psi(h_t, h_{t-1}) = y_{at} + b_t - (R_t + \mathcal{I}(b_{t-1} > 0) r_t^p) b_{t-1} + (1-\delta_H) h_{t-1} \tag{22}$$

$$\text{with} \quad b_t \le \min(m_{H,t} h_t, m_Y \mathfrak{R}_t), \quad c_t \ge 0, \ l_t \in (0, \bar{l}), \tag{23}$$

where $\mathcal{I}(b_{t-1} > 0)$ is the indicator function equal to 1 if the household is a net debtor, 0 otherwise. The state vector x_t remains unchanged with respect to the benchmark model, and so does the equilibrium definition. In the calibration, m_H drops by 6 percentage points in correspondence of the two lowest values of A_t, and remains constant for all other values of A_t.[22] The interest rate premium is 0.75% for the two lowest aggregate productivity realizations, in both periods (r^p is equal to zero for all other values of A_t). This simple modification can qualitatively account for the behavior of housing and debt in the 2007–2009 recession. Fig. 6 shows the impulse responses to positive and negative productivity shocks, comparing the early period with the late period (defined as in the baseline exercise). In the late period, debt, housing and GDP respond less to positive shocks, so that one finds evidence of the Great Moderation so long as the economy is lucky enough not to be hit by (too negative) negative shocks. When the worst recessionary shocks hit, however, the decline in debt and in housing purchases is much larger in the late period than in the early period. When debt is high, the housing sector can better absorb "small" business-cycle shocks, but becomes more vulnerable to large negative shocks that result in a credit crunch: these shocks cause highly leveraged households to sharply reduce their debt and housing purchases.

[21] Jermann and Quadrini (2012) document that credit shocks have played an important role in capturing U.S. output dynamics during the last decades.

[22] Total factor productivity is discretized using a 7-state Markov chain (see Appendix). For the lowest two aggregate productivity levels: in the period 1952–1982, $m_{H,t} = 0.70$, and in the period 1983–2010, $m_{H,t} = 0.80$.

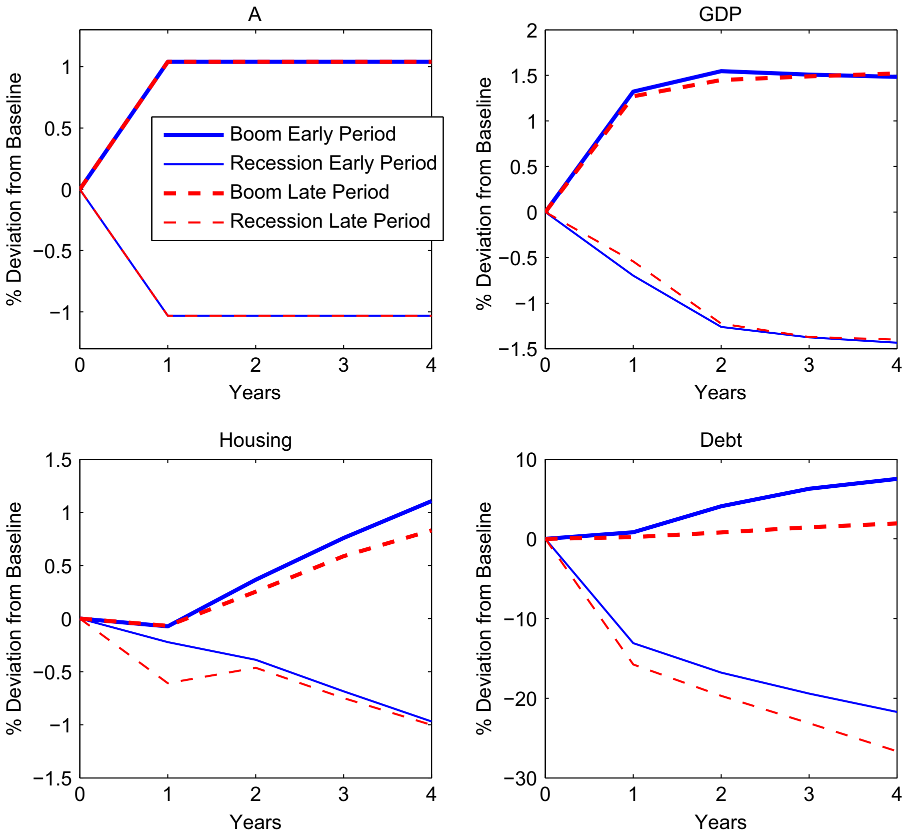

Fig. 6. Responses to positive and negative technology shocks: model with cyclical loan-to-value and interest rate premia. *Note*: Model dynamics following an exogenous switch in productivity *A* in period zero. The thick lines plot a 1% increase in productivity that does not change financial conditions in the early (solid lines) and late (dashed lines) period calibration. The thin lines plot a 1% decrease in productivity together with a worsening in financial conditions.

8. Sensitivity analysis

Table 5 presents the results for versions of the model with alternative calibrations.

Homogeneous discount factor. The discount factor ($\beta = 0.978$) and the utility from renting ($\theta = 0.922$) are chosen to achieve the same homeownership rate and interest rate as in the benchmark. The volatilities of housing investment and output are now slightly higher than in the baseline calibration, but the correlations of housing investment and hours with output fall: this result occurs since fewer people are close to the borrowing limit (only 15% of households are liquidity-constrained) and in need of increasing hours to finance the downpayment in good times. In addition, with a single discount factor, few people hold debt in equilibrium, and the distribution of wealth is more egalitarian than in the data: the Gini coefficient for wealth is 0.53, lower than in the data and in the benchmark model. The model predicts, unlike the data, a negative correlation between turnover and GDP: with a single discount rate, more housing capital reallocation occurs in bad times.

Income shocks persistence. Our benchmark sets $\rho_Z = 0.9$. Holding total income risk constant, some of the model properties are a non-monotonic function of ρ_Z. When the shocks are not very persistent ($\rho_Z = 0.7$), the equilibrium level of debt is lower, fewer people are liquidity-constrained, and debt and housing investment are less volatile and slightly less cyclical. When income shocks are highly persistent ($\rho_Z = 0.95$), more people are liquidity-constrained, but more people are lucky for a spell long enough to afford the downpayment for a house and to keep housing and debt relatively unchanged in response to shocks. In other experiments, only for intermediate values of the persistence coefficient (between 0.85 and 0.92), can the model account for both the high volatility of housing investment and the high correlation of debt with economic activity. Moreover, for values of ρ_Z above 0.95, housing turnover is negatively correlated with GDP, and housing is negatively correlated with business investment.

Housing transaction costs. The standard deviation of housing investment is 6.42% in the benchmark, and rises to 10.42% with zero transaction costs.[23] Because houses are less risky, homeownership rises, from 64% to 68%. Aggregate volatility

[23] Thomas (2002) argues that lumpiness of fixed investment at the level of a single production unit bears no implications for the behavior of aggregate quantities in an otherwise standard RBC model. Her argument rests on the representative household's desire to smooth consumption over

Table 5
Robustness analysis.

	Data	Model	One-β	Persistence		Transaction cost		Low δ
				$\rho_Z = 0.7$	$\rho_Z = 0.95$	$\psi = 0\%$	$\psi = 8\%$	$\delta_H = 3\%$
Standard deviation								
GDP	*2.09*	2.09	2.16	2.08	2.02	2.05	2.01	2.05
C	*0.93*	1.63	1.69	1.69	1.69	1.69	1.72	1.68
IH	*7.12*	6.42	6.72	4.99	4.73	10.42	3.45	11.33
IK	*4.90*	4.16	4.83	4.24	4.12	4.99	3.95	5.17
Debt	*2.23*	8.34	14.78	2.68	2.11	1.68	2.11	0.68
Hours	*1.60*	0.33	0.39	0.32	0.27	0.36	0.27	0.30
Housing turnover	*0.54*	0.29	0.40	0.16	0.22	2.14	0.13	0.16
Correlations								
IH, GDP	*0.89*	0.66	0.58	0.61	0.49	0.34	0.54	0.30
Debt, GDP	*0.78*	0.71	0.72	0.60	0.58	0.69	0.39	0.11
Hours, GDP	*0.82*	0.65	0.60	0.50	0.43	0.45	0.34	0.45
Turnover, GDP	*0.69*	0.39	−0.32	0.18	−0.15	0.67	−0.08	0.10
IH, IK	*0.36*	0.18	0.08	0.19	0.03	−0.40	0.19	−0.44
Debt, C	*0.72*	0.85	0.83	0.78	0.72	0.82	0.54	0.24
Averages								
Homeownership (%)	*64*	64	64	66	71	68	74	70
Debt to GDP (%)	*34*	31	9	17	42	40	37	46
Housing turnover (%)	*3.9*	4.0	3.3	4.7	2.9	42.0	2.1	3.8
Gini wealth	*0.79*	0.73	0.53	0.68	0.73	0.73	0.72	0.72
Gini labor income	*0.40*	0.41	0.42	0.45	0.39	0.41	0.41	0.42
Gini consumption	*0.23*	0.26	0.24	0.23	0.26	0.26	0.26	0.26
Liquidity constrained	*0.51*	0.45	0.15	0.30	0.49	0.47	0.45	0.45

Notes: In the one-β model, θ and average β are set so that the homeownership rate is 64% and the interest rate is 3%, as in the baseline model. No parameter changes are made in the other models, except those noted in row 2 of the table.

falls: housing and nonhousing capital are closer substitutes as means of saving, and the higher volatility of housing investment is offset by the reduced covariance between housing and nonhousing investment. The correlation between housing and non-housing investment, which is 0.18 in the benchmark (0.36 in the data), becomes −0.40 in absence of transaction costs. It is interesting to relate this result to the household production literature, which models adjustment costs either as convex or using a time-to-build specification. Fisher (2007) argues that the household production model predicts that housing and business investment are negatively correlated, unless one assumes that household capital is complementary to business capital and labor in market production. Our baseline model reproduces (unlike the model with no transaction costs) the positive empirical correlation between housing and business investment: sooner or later the transaction costs must be paid in order to consume more housing, and it is optimal to pay them in good times. Moreover, impatient renters cannot wait to become homeowners, thus effectively buying houses and borrowing (i.e. selling claims on capital) after a positive productivity shock.

The high adjustment cost case ($\psi = 8\%$) predicts low housing turnover (2.1%) relative to the data (4%), and an acyclical behavior of housing sales (sales are procyclical both in the data and in the benchmark model). Such model severely underpredicts the volatility of housing investment. Moving shocks (when combined with income shocks) could restore the level of housing turnover that is observed in the data even in the presence of high transaction costs: it is not clear, however, whether moving shocks could make turnover procyclical, unless they are more likely to happen in good times.

Housing depreciation. When the housing depreciation rate is lowered from 5% to 3%, the performance of some of the model's second moments worsens considerably. Housing investment becomes too volatile, the cyclicality of housing investment is much lower than in the data, and the model fails to match the comovement of housing with business investment.

9. Conclusions

The dynamics of housing and debt in our model are realistic not only at the macroeconomic level, but also at the level of individual household behavior: even if agents only infrequently adjust their housing choice, housing investment is the most volatile component of aggregate demand in our model, like in the data. Our model accounts for the procyclicality and

(*footnote continued*)
time, a desire that undoes any lumpiness at the level of the individual firm. Our sensitivity analysis shows that there are differences between the models with and without adjustment cost. Adjustment costs imply smaller housing adjustment at the aggregate level, but larger housing adjustments (when they occur) at the individual level.

volatility of housing investment, as well as for the procyclicality of household debt. The model can also explain why housing investment has become relatively less volatile, and household debt less procyclical, as a consequence of increased household-level risk and lower downpayment requirements, two structural changes that have occurred in the U.S. economy around the mid-1980s. We further extend the model to account for a "Great Recession" episode characterized by negative technology shocks coupled with tighter credit conditions. This simple modification generates an interesting nonlinearity which is consistent with recent events: when leverage is high, housing, debt and output respond less to positive shocks (as in the Great Moderation) but are relatively more vulnerable to negative shocks, making a recession worse (as in the Great Recession).

Despite its complexity, one limitation of our approach is that it focuses on housing investment and does not endogenize house prices. Allowing for variable house prices would require specifying a two-sector model with housing and nonhousing goods that are produced using different technologies, or a model with different price stickiness in housing and nonhousing goods; and would probably require a rich array of shocks in addition to productivity shocks, since we know from existing studies that technology shocks alone cannot quantitatively explain observed movements in house prices: all of this would considerably increase computational costs.

Our model does not explicitly consider default. When all debt is collateralized, if no shock is large enough to cause agents to owe on their house more than they are worth, agents never find it optimal to default. Appendix C sketches an extension that features large housing depreciation shocks as the source of business cycles. The model allows borrowers to default, at the cost of losing their house and being excluded from the mortgage market. Lenders cannot observe individual borrowers' characteristics, but can charge a higher interest rate on all loans in states of the world where default rates are higher to satisfy a zero profit condition. In this setup, indebted households weigh the utility premium benefit of being homeowners against the cost of servicing their debt in states where they have negative equity. When a depreciation shock destroys part of the housing capital, borrowing rates rise, highly leveraged individuals owe more than they own, and decide to default on their debt, becoming renters. The model can be used to study how shocks to housing values interact with the mortgage defaults, interest rates, debt and the housing stock. For plausibly calibrated values, a shock that destroys 20% of the housing stock leads to a rise in defaults (from 0% to 10%), a rise in borrowing premia (from 0% to 1.5%), and a sharp decline in debt, output and housing investment.

Acknowledgments

Thanks to Massimo Giovannini and Joachim Goeschel for their invaluable research assistance. Thanks to Chris Carroll, Geng Li, Kalin Nikolov, Dirk Krueger, Makoto Nakajima, and seminar and conference participants for helpful comments on various drafts of this paper. Pavan acknowledges financial support from the Spanish Ministry of Education (Programa de Movilidad de Jovenes Doctores Extranjeros) and from the Spanish Ministry of Science and Innovation (ECO2011-23634).

Appendix A. Supplementary data

Supplementary data associated with this article can be found in the online version at http://dx.doi.org/10.1016/j.jmoneco.2012.10.020.

References

Anily, S., Hornik, J., Israeli, M., 1999. Inferring the distribution of households' duration of residence from data on current residence time. Journal of Business & Economic Statistics 17, 373–381.

Bottazzi, R., Low, H., Wakefield, M., 2007. Why Do Home Owners Work Longer Hours? IFS Working Papers, Institute for Fiscal Studies.

Cagetti, M., 2003. Wealth accumulation over the life cycle and precautionary savings. Journal of Business & Economic Statistics 21, 339–353.

Campbell, J.R., Hercowitz, Z., 2005. The Role of Collateralized Household Debt in Macroeconomic Stabilization. NBER working papers, National Bureau of Economic Research, Inc.

Chambers, M., Garriga, C., Schlagenhauf, D.E., 2009. Housing policy and the progressivity of income taxation. Journal of Monetary Economics 56, 1116–1134.

Davis, M.A., Heathcote, J., 2005. Housing and the business cycle. International Economic Review 46, 751–784.

Davis, M.A., Ortalo-Magne, F., 2011. Household expenditures, wages, rents. Review of Economic Dynamics 14, 248–261.

Díaz, A., Luengo-Prado, M.J., 2010. The wealth distribution with durable goods. International Economic Review 51, 143–170.

Dynan, K.E., Elmendorf, D.W., Sichel, D.E., 2007. The evolution of household income volatility. In: Finance and Economics Discussion Series, Board of Governors of the Federal Reserve System, US.

Eberly, J.C., 1994. Adjustment of consumers' durables stocks: evidence from automobile purchases. Journal of Political Economy 102, 403–436.

Favilukis, J., Ludvigson, S., Nieuwerburgh, S.V., 2009. Macroeconomic Implications of Housing Wealth, Housing Finance, and Limited Risk-Sharing in General Equilibrium. Unpublished Manuscript.

Fernandez-Villaverde, J., Krueger, D., 2004. Consumption and saving over the life cycle: how important are consumer durables? In: 2004 Meeting Papers, Society for Economic Dynamics.

Fisher, J.D.M., 2007. Why does household investment lead business investment over the business cycle? Journal of Political Economy 115, 141–168.

Fisher, J.D.M., Gervais, M., 2007. First-Time Home Buyers and Residential Investment Volatility. Working Paper Series, Federal Reserve Bank of Chicago.

Floden, M., Lindé, J., 2001. Idiosyncratic risk in the United States and Sweden: is there a role for government insurance?. Review of Economic Dynamics 4, 406–437.

Foote, C., Hurst, E., Leahy, J., 2000. Testing the (s, s) model. American Economic Review 90, 116–119.

Fortin, N.M., 1995. Allocation inflexibilities, female labor supply, and housing assets accumulation: are women working to pay the mortgage?. Journal of Labor Economics 13, 524–557.
Francis, N., Ramey, V.A., 2009. Measures of per capita hours and their implications for the technology-hours debate. Journal of Money, Credit and Banking 41, 1071–1097.
Fraumeni, B., 1997. The measurement of depreciation in the U.S. national income and product accounts. Survey of Current Business, 7–23.
Gerardi, K.S., Rosen, H.S., Willen, P.S., 2010. The impact of deregulation and financial innovation on consumers: the case of the mortgage market. Journal of Finance 65, 333–360.
Gervais, M., 2002. Housing taxation and capital accumulation. Journal of Monetary Economics 49, 1461–1489.
Gomme, P., Kydland, F.E., Rupert, P., 2001. Home production meets time to build. Journal of Political Economy 109, 1115–1131.
Greenwood, J., Hercowitz, Z., 1991. The allocation of capital and time over the business cycle. Journal of Political Economy 99, 1188–1214.
Gross, D.B., Souleles, N.S., 2002. Do liquidity constraints and interest rates matter for consumer behavior? evidence from credit card data. The Quarterly Journal of Economics 117, 149–185.
Guvenen, F., 2011. Macroeconomics with Heterogeneity: A Practical Guide. Unpublished manuscript.
Hall, R.E., 2011. The long slump. The American Economic Review 101, 431–469.
Hansen, G.D., 1993. The cyclical and secular behaviour of the labour input: comparing efficiency units and hours worked. Journal of Applied Econometrics 8, 71–80.
Hansen, K.A., 1998. Seasonality of Moves and Duration of Residence. Current Population Reports, Household Economic Studies.
Hendricks, L., 2007. How important is discount rate heterogeneity for wealth inequality? Journal of Economic Dynamics and Control 31, 3042–3068.
Iacoviello, M., Neri, S., 2010. Housing market spillovers: evidence from an estimated dsge model. American Economic Journal: Macroeconomics 2, 125–164.
Jappelli, T., 1990. Who is credit constrained in the U.S. economy? The Quarterly Journal of Economics 105, 219–234.
Jermann, U., Quadrini, V., 2012. Macroeconomic effects of financial shocks. The American Economic Review 102, 238–271.
Johnson, K.W., Li, G., 2010. The debt-payment-to-income ratio as an indicator of borrowing constraints: evidence from two household surveys. Journal of Money, Credit and Banking 42, 1373–1390.
King, R.G., Rebelo, S.T., 1999. Resuscitating real business cycles. In: Taylor, J.B., Woodford, M. (Eds.), Handbook of Macroeconomics, 1st ed. vol. 1. Elsevier, pp. 927–1007 (Ch. 14).
Kiyotaki, N., Michaelides, A., Nikolov, K., 2011. Winners and losers in housing markets. Journal of Money, Credit and Banking 43, 255–296.
Krueger, D., Perri, F., 2006. Does income inequality lead to consumption inequality? Evidence and theory. Review of Economic Studies 73, 163–193.
Krusell, P., Smith, A.A., 1998. Income and wealth heterogeneity in the macroeconomy. Journal of Political Economy 106, 867–896.
Moffitt, R., Gottschalk, P., 2008. Trends in the Transitory Variance of Male Earnings in the U.S., 1970–2004. Boston College Working Papers in Economics.
National Association of Realtors, 2005. Structure, Conduct, and Performance of the Real Estate Brokerage Industry. Technical Report.
Silos, P., 2007. Housing, portfolio choice and the macroeconomy. Journal of Economic Dynamics and Control 31, 2774–2801.
Tauchen, G., 1986. Finite state Markov-chain approximations to univariate and vector autoregressions. Economics Letters 20, 177–181.
Thomas, J.K., 2002. Is lumpy investment relevant for the business cycle? Journal of Political Economy 110, 508–534.
Wolff, E.N., 2010. Recent Trends in Household Wealth in the United States—Rising Debt and the Middle-Class Squeeze—An Update to 2007. Economics Working Paper Archive, Levy Economics Institute.
Yoshikawa, H., Ohtaka, F., 1989. An analysis of female labor supply, housing demand and the saving rate in Japan. European Economic Review 33, 997–1023.
Young, E.R., 2010. Solving the incomplete markets model with aggregate uncertainty using the Krusell–Smith algorithm and non-stochastic simulations. Journal of Economic Dynamics and Control 34, 36–41.
Ziliak, J.P., Kniesner, T.J., 1999. Estimating life cycle labor supply tax effects. Journal of Political Economy 107, 326–359.