Housing and Debt over the Life Cycle and over the Business Cycle*

Matteo Iacoviello†
Boston College

Marina Pavan‡
The Geary Institute, University College Dublin

November 24, 2008

VERY PRELIMINARY DRAFT. COMMENTS WELCOME.

Abstract

We present an equilibrium life-cycle model of housing tenure where nonconvex adjustment costs lead individuals to adjust housing infrequently and by big amounts. In the cross-sectional dimension, the model matches the wealth distribution, the age profiles of consumption, homeownership and mortgage debt, and data on the frequency of housing adjustment. In the time series dimension, the model accounts for the procyclicality and volatility of housing investment, and for the procyclical behavior of household debt.

We use a calibrated version of our model to ask: what are consequences for aggregate volatility of an increase in individual income risk and a decrease in down-payment requirements? We distinguish between an early period (1950s through 1970s), when individual risk was relatively small and loan-to-value ratios were low; and a late period (1980s through today), with high individual risk and high loan-to-value ratios. In the early period, precautionary saving is small, wealth-poor people are close to the maximum borrowing limit, and housing investment, homeowership and debt closely track aggregate productivity. In the late period, precautionary saving is larger, wealth-poor people borrow less than the maximum and become more cautious in response to aggregate shocks. As a consequence, the correlation between debt and economic activity on the one hand, and the sensitivity of housing investment to aggregate shocks on the other, are lower, as in the data. Quantitatively, our model can explain: (1) 30 percent of the reduction in the volatility of household investment; (2) the sharp decline in the correlation between household debt and economic activity; (3) between 5 and 10 percent of the reduction in the volatility of GDP.

*We are grateful to Massimo Giovannini and Joachim Goeschel for their invaluable research assistance. We thank Kalin Nikolov as well seminar participants at the Boston Fed, the Society for Economic Dynamics 2008 conference, the Bank of England, the ECB, Georgetown University, the Richmond Fed and the University of Amsterdam for helpful comments.

†iacoviel@bc.edu. Address: Boston College, Department of Economics, 140 Commonwealth Ave, Chestnut Hill, MA, 02467, USA.

‡marina.pavan@ucd.ie. Address: Geary Institute, University College Dublin, Belfield, Dublin 4, Ireland.

1. Introduction

The Setup. Housing investment is a volatile component of GDP. Historically, this observation has led researchers to emphasize movements in the housing market as central to understanding aggregate fluctuations. Despite this, modern business cycle theory has been surprisingly silent on this topic. When housing is included in equilibrium models,[1] its role is inconsistent with its definition: there is no role for income and wealth hetereogeneity, no borrowing constraints, no distinction between owning and renting, no transaction costs (or irrealistic ones) for adjusting home size, no life-cycle considerations.

Our goal in this paper is to address this imbalance. Namely, we study the business cycle and the life-cycle properties of household investment and household debt in a quantitative general equilibrium model. To this end, we modify a standard life-cycle model (in which households face idiosyncratic income and mortality risk) to allow for aggregate uncertainty, on the one hand, and for an explicit treatment of housing, on the other. We introduce aggregate uncertainty by making aggregate total factor productivity time-varying. We introduce housing by carefully modeling some key features that make housing different from other goods: its role as a collateral for loans, its lumpiness, and the choice of renting vs. owning. Finally, we relax the assumption that individuals have identical tastes by splitting the population in a patient group and in an impatient group: this simple modification makes the wealth distribution highly skewed, in a manner similar to the data.[2]

The Results. When calibrated to reproduce some key observations of post-world-war II U.S. economy, our model does a good job in accounting for several facts.

At the cross-sectional level, our model reproduces the U.S. wealth distribution almost perfectly, and replicates well the life-cycle profile of housing and non-housing wealth. The very young and the very old (as well as the poor) are renters, and hold very few assets. The middle-age and the wealth-rich become homeowners. For a typical household, the portfolio of assets is remarkably simple: a house, and a large mortgage. Despite its stylized nature, the model also reproduces frequency and size of microeconomic housing adjustment: homeowners change house size often and by big amounts; renters change house size often, but by smaller amounts.

In terms of its business cycle properties, our model successfully replicates two empirical properties of housing investment: (1) its procyclicality, and (2) its high volatility. In addition, the model is able to match (3) the procyclical behavior of household mortgage debt. To our knowledge, no previous model with rigorous micro-foundations for housing demand has succeeded in reproducing these regularities in quantitative general equilibrium.

[1] See Davis and Heathcote (2005) and Fisher (2007) for examples of this kind.

[2] Krusell and Smith (1998) explore a heterogeneous-agents model without housing and with discount rate heterogeneity which replicates the observed data on the distribution of wealth.

The Model Experiments. We illustrate the workings of our model with two experiments, by characterizing the business cycle implications of (1) increasing microeconomic volatility (idiosyncratic risk) and (2) lowering down-payment constraints. Both these structural changes might have affected the sensitivity of macroeconomic aggregates to given economic shocks, and are potential candidates in explaining the role of debt and the housing market in the Great Moderation, especially given two observations on the U.S. economy post 1980's (see Figure 1 and Table 1): first, the volatility of housing investment has fallen more than proportionally relative to GDP; second, the correlation between mortgage debt and economic activity has dropped, from 0.78 to 0.16.[3]

In line with the data, we find that the combination of larger idiosyncratic risk and lower down-payment requirements can (1) reduce the relative volatility of housing investment; (2) reduce the correlation between household debt and GDP.

Lower downpayment requirements increase homeownership rates, by making it easier to buy a house. The higher number of homeowners changes the business cycle properties of the economy for two reasons. First, indebted homeowners are more likely to work more in bad times (relative to renters) in order to finance housing and mortgage payments, thus offsetting the decrease in output due to negative TFP shocks. Second, homeowners are also less likely to adjust their housing capital over the business cycle (compared to an economy with a higher number of renters who can become first time buyers). Both these forces reduce housing investment volatility and aggregate volatility.

An increase in income risk per se leads to higher precautionary saving, and to a decrease in homeownership rates among impatient individuals. In addition, larger risk makes individuals more cautious, in the sense that they react less to aggregate productivity shocks. Combined with low downpayment requirements, this effect reduces the procyclicality of household debt.

We find that only together the change in income volatility and in downpayment requirements go in the direction of explaining the changes observed in the data. Taken together, these two mechanisms can explain 30 percent of the reduction in the variance of housing investment, and the entire decline in the correlation between debt and economic activity.

Previous Literature. Our model is part of a large and growing literature that analyzes the aggregate behavior of economies with heterogeneous agents, incomplete markets and aggregate shocks.[4] However, most of this literature abstracts from housing altogether, or implicitly considers housing as part of the economywide capital stock.[5] Some exceptions are listed below.

[3] See for instance Campbell and Hercowitz (2006) for the role of financial reforms, and Dynan, Elmendorf and Sichel (2007) for a discussion on the evolution of household income volatility.

[4] Most of these model economies feature an "approximate aggregation" result (see Krusell and Smith, 1998): that is, their business cycle properties can be described with reference to a small set of aggregate variables, such as the mean of the wealth distribution.

[5] Papers on housing in incomplete market models with heterogeneous agents that abstract from aggregate shocks include Gervais (2002), Fernandez-Villaverde and Krueger (2005), Nakajima (2005), Silos (2005), Diaz and Luengo-Prado (2005) and Ortalo-Magne' and Rady (2006) and Gruber and Martin (2006) among others.

Silos (2007) analyzes the relationship between macroeconomic shocks and household portfolio choice adopting a life-cycle framework, but he does not model the extensive margin of owning versus renting and assumes irrealistic convex costs for housing adjustment.[6] His focus is on the impact of aggregate shocks on the wealth distribution and portfolio composition. On the opposite, we concentrate on how individual risk and different downpayment requirements affect macroeconomic fluctuations through household portfolio choices.

Other papers analyze housing or durable goods in the context of equilibrium business cycle models which share some features with ours. Fisher and Gervais (2007) find that the decline in residential investment volatility is driven by a change in the demographics of the population together with an increase in the cross-sectional variance of earnings. However, theirs is in a certain sense a partial equilibrium framework in which the interest rate does not react to aggregate fluctuations. Campbell and Hercowitz (2006) study the impact of financial innovation on macroeconomic volatility, and their mechanism is through the labor supply: less tight collateral constraints weaken the connection between constrained households' housing investment and their hours worked. Kiyotaki, Michealides and Nikolov (2007) use a stylized life-cycle model to study the interaction between borrowing constraints, house and land prices, and economic activity.

Finally, our modeling approach shares some features with papers that, abstracting from housing, have analyzed business cycle fluctuations in life-cycle economies. Notable examples in this literature include Rios-Rull (1996) and Gomme et al. (2004).

2. The Model Economy

Our benchmark economy is a version of the stochastic growth model with overlapping generations of heterogeneous consumers, extended to allow for housing investment, collateralized borrowing and a housing rental market.

Time is discrete. Individuals live at most T periods and work until age $\widetilde{T} < T$. Agents' labor endowment depends on a deterministic age-specific productivity and a stochastic component, whose process is exogenously specified. Retirement is mandatory and people receive a lump-sum pension P every period starting at age $\widetilde{T}+1$. When an agent dies, a new agent of the same dinasty is born, with earnings that are correlated with the dead person's earnings. Denote χ_{a+1} the probability of surviving from age a to $a+1$, and let Π_a be the stationary distribution over ages. Each period a generation is born of the same measure of dead agents, so that the total measure of individuals does not change over time: $\Pi_1 = \sum_{a=1}^{T}(1-\chi_{a+1})$. Let the measure of all individuals at any given period

[6] Convex adjustment costs for housing induce adjustment dynamics different from the specification of adjustment costs we use in this paper. Under the convex specification, housing adjustment takes the form of a series of small adjustments over a number of periods. Under our specification, the homeowner's housing stock follows an (S,s) rule, remaining unchanged over a long period and ultimately changing by a potentially large amount. Modeling the adjustment cost as proportional to the stock seems much more plausible for housing.

be normalized to one: $\sum_{a=1}^{T} \Pi_a = 1$.

At each point in time, agents may differ in three respects:

1. Their age;

2. Their realized labor productivity;

3. Their degree of patience: a recent branch of this literature suggests that preference heterogeneity may be an important source of wealth inequality. This is motivated by the finding that similar households hold very different amounts of wealth. For example, Venti and Wise (2001) study wealth inequality at the outset of retirement among households with similar lifetime earnings and conclude that "the bulk of the dispersion must be attributed to differences in the amount that households choose to save".[7]

Consumers have preferences over a non-durable consumption good, housing services and leisure. They can choose between renting housing services in a rental market and owning housing capital.

There are no state contingent markets for hedging against idiosyncratic risk, and only self-insurance through a risk-free bond is possible. Agents can borrow up to a fraction of their housing wealth, and incur a cost in adjusting the housing stock. Finally, aggregate uncertainty is introduced in the form of a shock to total factor productivity. Hence the model uses as inputs the exogenous aggregate and idiosyncratic uncertainty, and delivers as output the endogenously derived dynamics of housing and financial investments over the life cycle and the business cycle.

2.1. Household Preferences

Let $\bar{l}$ denote each agent's total time endowment. Households derive utility from leisure $(\bar{l} - l)$, non-durable consumption c, and service flows s from housing, which are assumed to be proportional to the housing stock owned or rented.

The per-period utility function is additively separable in its arguments, and takes the following formulation:

$$u\left(c, s, \bar{l} - l\right) = \frac{[c(\theta s)^{j} \left(\bar{l} - l\right)^{\tau}]^{1-\sigma} - 1}{1 - \sigma}$$

with $j, \tau > 0$, the risk aversion coefficient σ is positive, and $\theta = 1$ if $s = h > 0$ (the individual owns his house), $\theta < 1$ if $h = 0$ (the individual rents). The assumption on θ implies that a household experiences a net utility gain when transiting from renting to owning a home and is standard in models of home ownership in the public economics and urban economics literature, like Rosen (1985) and Poterba (1992), among others. We also assume that, whenever individuals own, there is a

[7]Krusell and Smith (1998) also explore a heterogeneous-agents setting with preference heterogeneity which, unlike a benchmark model with a single discount factor, replicates some of the key features of the observed data on the distribution of wealth.

minimum size house that can be purchased, so that rental units may come in smaller sizes than houses, thus allowing renters to consume smaller amount of housing services, as in Gervais (2002).

Each unit of time supplied in the labor market in period t is paid at the wage rate w_t. The total 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, $z \in \widetilde{Z} \equiv \left\{z^1, ..., z^n\right\}$. The shock follows a Markov process with transition matrix $\pi_{z,z'} = \Pr\left(z_{t+1} = z' | z_t = z\right)$, $\pi_{z,z'} > 0$ for every $z, z' \in \widetilde{Z}$, with $\sum_{z'} \pi_{z,z'} = 1$ for every $z \in \widetilde{Z}$. By a law of large numbers, π also represents the fraction of agents experiencing a transition from z to z' between any two periods, with z and $z' \in \widetilde{Z}$. Let Π be the unique stationary distribution associated with the transition probability π. Again, by a law of large numbers, at each period there are $\Pi(z)$ agents characterized by labor productivity z. The total amount of labor efficiency units $\sum_{i=1}^{n} z^i \Pi\left(z^i\right)$ as well as the sum of age-specific productivity values $\sum_{a=1}^{\widetilde{T}} \eta_a \Pi_a$ are both constant and normalized to one for convenience. From age $\widetilde{T}+1$ 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. Denote y_{at} the total net income at age a in period t. Then

$$
\begin{aligned}
y_{at} &= w_t \eta_a z_t l_t - \Gamma \text{ if } a \leq \widetilde{T} \\
y_{at} &= P \text{ if } a > \widetilde{T}.
\end{aligned}
$$

Households can buy and sell only one bond, b, which pays a gross risk-free interest rate of R_t in period t. For convenience, let positive amounts of this bond denote a net debt position.[8] Housing wealth can be used as collateral for borrowing. At any period, households can borrow up to a fraction $m_h < 1$ of their housing stock and a fraction m_y of a measure of their expected lifetime earnings:

$$b_t \leq \min\{m_h h_t, m_y \Re_t\left(y_{at}; R_t, w_t\right)\}$$

where $\Re_t\left(y_{at}; R_t, w_t\right) = y_{at} + \sum_{s=a+1}^{T} \frac{E_t(y_s | y_{at}; w_t)}{(R_t)^{s-a}}$ is computed at the current wage and interest rate, and is meant to capture the approximated present discounted value of one's lifetime labor earnings (and pension).[9]

Essentially, the specification of borrowing constraint (1) rules out unsecured debt; (2) restricts debt to homeowners only; (3) implies that the collateral constraint on h_t is more likely to bind early in life, when the present discounted value of earnings is high, while the constraint on $\Re_t$ is more likely to bind late in life, when the present discounted value of earnings is low.

Given that there are no state contingent markets for the individual shocks, the agent is able to smooth consumption only by adjusting the level of financial stock and housing stock over time. Each

[8] We therefore refer to b as financial liabilities (or net debt), and correspondingly to $-b$ as financial assets (or net assets).

[9] The measure is approximated since the interest rate is fixed at the current value. This greatly simplifies the calculation of the expected lifetime earnings, allowing us to derive a constraint which prevents the elderly from borrowing too much.

agent starts out with initial assets (b_0, h_0), the bequest left by a dead agent. Bequests left by people who die before T are interpreted as merely accidental. At the last age, the agent may derive utility from the bequest of his net wealth $h_T - b_T$:

$$u^b (h_T - b_T) = \omega^b \frac{(h_T - b_T)^{1-\sigma} - 1}{1 - \sigma}.$$

As a final twist, we introduce a life-cycle preference shifter λ_a in the utility function. Changes in λ_a capture in a crude way changes in household size that deterministically affect the marginal utility of consumption, as in Cagetti (2003).[10]

Summing up, agents maximize their expected lifetime utility:

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

where β_i is the household specific discount factor, $\beta_i \in (0, 1)$, and E_1 denotes expectations at age $a = 1$. We refer to households with a lower value of β as impatient. The discount factor is deterministic, and does not vary over time. In the numerical experiments below, we assume that households are either born impatient or patient. Heterogeneity in discount factors allows the model matching the wealth distribution better than in a model without these features.

2.2. The Financial Sector and the Housing Rental Market

A perfectly competitive intermediary financial sector collects deposits from households who save, lends to firms and households who borrow and to firms, and buys residential capital to be rented to households who choose to become tenants. We assume that the financial sector can convert each unit of the final good into one unit of physical capital (residential or not) without incurring any cost.[11] Let p_t^s be the price of each unit of rental services at time t. Then a no-arbitrage condition holds such that the net revenue from lending one unit of financial capital must be equal to the net revenue from renting one unit of housing capital:

$$p_t^s = E_t \left(\frac{R_{t+1} - (1 - \delta_h)}{R_{t+1}} \right)$$

at any t,[12] where δ_H is the depreciation rate of the housing stock.[13]

[10] We calibrate λ_a with a polynomial function which is taken from Cagetti (2003). It shows a hump shape over the life-cycle and a peak around age 40, thus reinforcing the hump in the life-cycle consumption profile.

[11] The financial sector operates the technology to transform output into capital, by purchasing output from the household sector and then investing into new capital, while earning revenues by renting existing capital to production firms. The financial sector finances itself by issuing deposits in t that pay a gross interest rate R_{t+1} in period $t + 1$.

[12] Alternatively, one can interpret the marginal cost of one house to be 1 for the financial sector, since output can be converted into housing costlessly; and the marginal benefit to be the sum of the current rental income, p_t^s, plus expected return next period, $E_t \left(\frac{1-\delta_h}{R_{t+1}} \right)$, where R_t is the opportunity cost of funds for the financial sector. Equating costs and benefits yields the equation in the text.

[13] The expectation term in the no-arbitrage condition is an outcome of the discrete time assumption: namely, it reflects the assumption that housing services are paid for and yield utility in the current period, whereas capital services are

2.3. Production

The goods market is perfectly competitive and characterized by constant returns to scale, so that without loss of generality we can consider a single, representative firm only. The good is produced according to the Cobb-Douglas technology:

$$Y = AK^{\alpha}L^{1-\alpha}$$

where L and K denote aggregate labor and aggregate capital respectively, $\alpha \in (0,1)$ is the capital share of aggregate income, while $A \in \widetilde{A} \equiv \left\{A^1, .., A^a\right\}$ represents the stochastic shock to total factor productivity. The aggregate shock is assumed to follow a finite-state Markov process with transition matrix $\pi_{A,A'} = \Pr\left(A_{t+1} = A' | A_t = A\right)$, with $\pi_{A,A'} > 0$ for every $A, A' \in \widetilde{A}$, and $\sum_{A'} \pi_{A,A'} = 1$ for every $A \in \widetilde{A}$.

The economy-wide feasibility constraint requires that at each period t total production of the good, Y_t, corresponds to 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 (owned and rented), and the total transaction costs incurred by homeowners for adjustments to the housing stock, which we denote by Ω:

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

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

We assume that there is no supply or demand of bonds from abroad. Hence the net supply of financial assets in this economy must be equal to the aggregate level of physical capital K_t plus the rented residential capital H_t^r. Factor prices will be determined in equilibrium by the optimization conditions of the representative firm, which maximizes its profits.

2.4. The Household Problem and Equilibrium

Denote with $\Phi_t\left(z_t, b_{t-1}, h_{t-1}; \beta, a\right)$ the distribution over productivity shocks, asset holdings, housing wealth (owned by the household), discount factors and ages in period t.[14] Without aggregate uncertainty, the economy would be in a stationary equilibrium, with an invariant distribution Φ and constant prices. Given aggregate volatility however, the distribution Φ will change over time, depending on the evolution of aggregate shocks and the heterogeneity of individual states at any period.

When solving their dynamic optimization problem, agents need to predict the future wage and interest rate. The latter depend on the future productivity shock and aggregate capital-labor ratio,

paid for and yield output and with a one-period delay (this timing assumption is standard in decentralized version of the neoclassical growth model, see for instance Ljungqvist and Sargent, 2004, chapter 12).

[14] The marginal distribution of Φ_t with respect to z_t is Π; that with respect to a is Π_a, by definition.

which in turn is determined by the overall distribution of individual states.[15] As a consequence, the distribution $\Phi_t(z_t, b_{t-1}, h_{t-1}; \beta, a)$ (and its law of motion) is one of the aggregate state variables that agents need to know in order to make their decisions (together with total factor productivity). This distribution is an infinite-dimensional object, and its law of motion maps an infinite-dimensional space into itself, which imposes a crucial complication for the solution of the model economy. Indeed, it is impossible to directly compute the equilibrium for such an economy. We thus adopt the computational strategy of Krusell and Smith (1998) and assume that only the first moments of the distribution Φ are sufficient to forecast future prices. Krusell and Smith's approach can be seen as a mere computational device to solve for an "approximate" equilibrium in this kind of models. In a different interpretation, agents could be thought of as having "partial information", or being characterized by "bounded rationality". In any case, Krusell and Smith (1998) show that their methodology is accurate enough so to have very small forecasting errors and an "approximate" equilibrium that is very close to the exact one.

We can write the household problem in recursive formulation. The agent's individual state variables are the productivity shock z_t, the net liabilities position b_{t-1}, and the stock of housing wealth h_{t-1} owned at the beginning of period t. In the spirit of Krusell and Smith (1998), agents need only forecast the future ratio of aggregate capital to aggregate labor in order to predict the next period' wage and interest rate. They observe aggregate capital K_{t-1}/L_t at the beginning of period t, and approximate the evolution of this variable with a linear function that depends on the aggregate shock A_t. Denote $x_t \equiv (z_t, b_{t-1}, h_{t-1}, A_t, K_{t-1}/L_t)$ the vector of individual and aggregate state variables.[16] In recursive form, the dynamic problem of an age a household with discount factor β_i can be stated as follows:

$$V_a(x_t; \beta_i) = \max_{I^h \in \{0,1\}} \{I^h V_a^h(x_t; \beta_i) + \left(1 - I^h\right) V_a^r(x_t; \beta_i)\} \tag{2}$$

where V_a^h and V_a^r are the value functions if the agent owns and rents a house, respectively, and $I^h = 1$ corresponds to the decision of owning. The value of being a homeowner is given by the following:

$$V_a^h(x_t; \beta_i) = \max_{c_t, b_t, h_t, l_t} \{\lambda_a u\left(c_t, h_t, \bar{l} - l_t\right) + \beta_i \chi_{a+1} \textstyle\sum_{z_{t+1}, A_{t+1}} \pi_{A_t, A_{t+1}} \pi_{z_t, z_{t+1}} V_{a+1}(x_{t+1}; \beta_i)\} \tag{3}$$

$$\text{s.t.} \qquad c_t + h_t + \Psi(h_t, h_{t-1}) = y_{at} + b_t - R_t b_{t-1} + (1 - \delta_H) h_{t-1}$$

$$b_t \leq \min\{m_h h_t, m_y \Re_t\}, \; c_t \geqslant 0, \; l_t \in \left(0, \bar{l}\right)$$

$$K_t/L_{t+1} = F(K_{t-1}/L_t; A_t)$$

[15]In an economy with heterogeneous agents, people have different incomes and different propensities to save out of their wealth, so that the entire wealth distribution (together with the optimal policy functions) is needed to predict the future aggregate capital-labor ratio, and thus the interest rate such that firms' profits are maximized.

[16]We present the "approximate" recursive formulation in which the aggregate state variables are represented by the economy' capital-labor ratio and the aggregate shock. As described in the text, in the "true" definition of the household dynamic problem the entire distribution Φ_t is an argument of the value function.

where $\digamma$ is a linear function in K_{t-1}/L_t, whose parameters depend on the aggregate shock A_t, and denotes the law of motion of the aggregate state, which agents take as given. The last term on the left-hand side of the budget constraint shows that measures the cost, proportional to his initial housing stock, that an owner has to pay 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 practice in the housing market that requires, for instance, fees paid to realtors to be equal to a fraction of the value of the house that is being sold.[17]

The value of renting a house is determined by solving the problem:

$$V_a^r(x_t;\beta_i) = \max_{c_t,b_t,s_t,l_t} \left\{\lambda_a u\left(c_t, s_t, \bar{l} - l_t\right) + \beta_i \chi_{a+1} \textstyle\sum_{z_{t+1},A_{t+1}} \pi_{A_t,A_{t+1}} \pi_{z_t,z_{t+1}} V_{a+1}(x_{t+1};\beta_i)\right\} \tag{4}$$

$$\begin{aligned} \text{s.t.} \quad & c_t + p_t^s s_t + \Psi(0, h_{t-1}) = y_{at} + b_t - R_t b_{t-1} + (1-\delta_H) h_{t-1} \\ & b_t \le 0,\ c_t \geqslant 0,\ l_t \in \left(0, \bar{l}\right), h_t = 0 \\ & K_t/L_{t+1} = \digamma\left(K_{t-1}/L_t; A_t\right) \end{aligned}$$

At the last age, $V_{T+1}(x_{T+1};\beta) = u^b(h_T - b_T)$.

We are now ready to formally define the equilibrium for this economy.

Definition 2.1. *A recursive competitive equilibrium is value functions* $\{V_a(x_t;\beta)\}_{a=1,..,T;t=1,..,\infty}$, *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 and period* t, *prices* $\{R_t\}_{t=1}^{\infty}$, $\{w_t\}_{t=1}^{\infty}$ *and* $\{p_t^s\}_{t=1}^{\infty}$, *aggregate variables* K_t, L_t, H_t *and* H_t^r *for each period* t, *lump-sum taxes* Γ *and pension* P, *and a law of motion* $\digamma(K_{t-1}/L_t; A)$ *such that:*

1. ***Agents optimize:*** *Given* R_t, w_t, p_t^s, *and the law of motion* $\digamma$, *the value functions are solution to the individual's problem, with the corresponding policy functions;*

2. ***Factor prices are determined competitively at any*** t***:***

$$R_t - 1 + \delta_K = \alpha A_t \left(L_t/K_{t-1}\right)^{(1-\alpha)}$$

$$w_t = (1-\alpha) A_t (K_{t-1}/L_t)^{\alpha}$$

and ***the rental price is given by the no-arbitrage condition****:*

$$p_t^s = E_t\left(\frac{R_{t+1} - (1-\delta_h)}{R_{t+1}}\right).$$

[17] Implicit in this formulation is the simplifying assumption the household can adjust the level of housing consumption only by selling the old house and buying thr new one. In practice, some adjustment to the level of housing consumption can be accomplished at the intensive margin while staying in the current house, to the extent that the household can expand, remodel, or fail to maintain the house. For simplicity, we rule this possibility out here.

3. ***The assets market clears at any*** t:

$$-\int b_a\left(x_t;\beta\right)\partial\Phi_t\left(z_t,b_{t-1},h_{t-1};\beta,a\right)=K_t+H_t^r,$$

where

$$H_t^r=\int(1-I_a^h\left(x_t;\beta\right))s_a\left(x_t;\beta\right)\partial\Phi_t\left(z_t,b_{t-1},h_{t-1};\beta,a\right).$$

4. ***The labor market clears at any*** t:

$$L_t=\int l_a\left(x_t;\beta\right)\eta_a z_t\partial\Phi_t\left(z_t,b_{t-1},h_{t-1};\beta,a\right)$$

and as a consequence the goods market satisfies the resource feasibility constraint (1), where H_t *is defined as follows:*

$$H_t=\int I_a^h\left(x_t;\beta\right)h_a\left(x_t;\beta\right)\partial\Phi_t\left(z_t,b_{t-1},h_{t-1};\beta,a\right)+H_t^r.$$

5. ***The Government budget is balanced***:

$$\sum_{a=1}^{\widetilde{T}}\Pi_a\Gamma=\sum_{a=\widetilde{T}+1}^{T}\Pi_a P.$$

6. ***The law of motion for the aggregate capital over labor ratio*** *is given by*

$$K_t/L_{t+1}=\Gamma\left(K_{t-1}/L_t;A_t\right).$$

Appendix A at the end provides the details on the computational approach used to solve for the model equilibrium.

3. Parameterization

In our baseline calibration, we aim at reproducing basic facts for the U.S. economy from 1952 to 1982. We characterize this period as having high aggregate volatility, low idiosyncratic volatility, and high downpayment requirements. Later, we will change idiosyncratic volatility and downpayment requirements in order to see the role that they play in affecting the properties of economic aggregates.

3.1. Demographics

One period in our model is a year. We assume that the economically active life of a household starts at age 21. Agents work $\widetilde{T}=45$ years until they reach age 65. They live off their savings and a lump-sum pension thereafter. Each period, the sequence of conditional survival probabilities is set equal to the survival probabilities for men aged 21-90, taken from the U.S. Decennial Life Tables for

1989-91.[18] We truncate the distribution at age 90, so that agents die with certainty on their 91st birthday. Each period, the measure of those who are born is equal to the measure of those who die. As a result, total population is constant.

The age polynomial λ_a, which captures the effect of demographic variables in the utility function, is taken from Cagetti (2003, figure B.2 in his paper) and approximated using a fourth-order polynominal. After normalizing the household size to unity at age 21, the household size peaks at 2.5 at age 40, and declines slowly to about 1 around age 90.

3.2. Endowments

We take the deterministic age-profile of efficiency units of labor for males aged 21-65 from Hansen (1993) and approximate it using a quadratic polynomial. The ratio of peak productivity to productivity at age of 21 is 1.8 and occurs at age 50. We impose mandatory retirement at age 65. Upon retirement and until death, each agent receives a pension equal to 43 (50) percent of the average gross (net) labor income in the economy, which is financed through lump-sum taxes.

Several studies document the increase in the cross-sectional dispersion of earnings in the United States between the 1970's and the 1990's. This increase is often decomposed into a rise in permanent inequality (attributable to education, experience, sex, etc.) and a rise of the persistent or transitory shocks volatility. Despite some disagreement on the relative importance of these two components, the literature finds that both play a role in explaining the increase in income dispersion (see Appendix B for details).

For our purposes, we are interested in the role played by the volatility of the persistent shocks to earnings. Therefore our stochastic idiosyncratic shock to labor productivity is specified as an autoregressive process of order one as follows:

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

which we approximate with a three-state Markov process.[19] We set $\rho_Z = 0.9$. In the baseline calibration we use $\sigma_Z = 0.30$, and we increase this number to $\sigma_Z = 0.45$ to capture the increased earnings volatility of the 1990s. Appendix B discusses the calibration of these numbers, which are in the ballpark of the microeconometric estimates, in more detail.

3.3. Preferences and Housing Costs

We relax the assumption that all consumers have identical tastes. Specifically, we assume that there are two classes of consumers, a 'patient' group with a discount factor of 0.995 and an 'impatient' group with a a discount factor of 0.925. We further assume that the impatient consumers compose 2/3 of the population. The high discount factor pins the average real interest rate down to an

[18] See http://www.cdc.gov/nchs/data/lifetables/life89_1_1.pdf

[19] We use Tauchen (1986) procedure.

average value around 3 percent, as in the data. The low discount factor is in the range of estimates in the literature, see for instance Hendricks (2007) and references therein. The shares of patient and impatient agents imply that 1/3 of people hold most of wealth, and deliver a Gini coefficient for wealth around 0.75, in line with the data.

Further, we assume logarithmic preferences, so that $\sigma = 1$. We set $\tau = 1.5$ and the total endowment of time $\bar{l} = 2.5$: these two parameters imply that the time spent working in the market (around 1.1) is around 43 percent of the available time.

We set the weight on housing in the utility function $j = 0.15$, and the depreciation rate for housing $\delta_H = 0.05$. These parameters yield average housing stock to output ratios of around 1.3 and average housing investment to output ratios around 6% on an annual basis. These values are in accordance with the National Income and Product Accounts and the Fixed Assets Tables.

The household incurs a proportional cost equal to $\psi = 4\%$ of the current housing stock if its net housing investment changes. We interpret this cost as a moving cost.

3.4. Technology

We set the capital share $\alpha = 0.33$ and its depreciation rate $\delta_K = 0.10$. In all the economies we consider, these values yield average capital to output ratios around 2.6 and average business investment to output ratios around 25% on an annual basis.[20]

Our calibration of the aggregate shock is meant to reproduce a standard deviation of output that matches the data counterpart for the period 1952-1982. We use a Markov-chain specification for aggregate productivity with five states to match the following first-order autoregressive representation for the logarithm of total factor productivity:

$$\log A_t = \rho_A \log A_{t-1} + \sigma \left(1 - \rho_A^2\right)^{1/2} \varepsilon_t, \qquad \varepsilon_t \sim Normal\left(0, 1\right).$$

We set $\rho_A = 0.90$ and $\sigma_A = 0.0152$. The first number mimics (after rounding) a quarterly autocorrelation rate of productivity of 0.979 as reported in King and Rebelo (1999). The second number is calibrated to match the standard deviation of HP-filtered output to that of the data.

Last, we set the maximum loan-to-value ratio m_h at 0.75 in the baseline calibration. We increase this number to 0.85 in the late period. The value of m_y is set equal to 0.25 in the baseline calibration and to 0.5 in the calibration for the late period: the income constraint is binding late in life, and essentially prevents old homeowners from borrowing too much late in life. Aside from this, our chosen value for m_y is large enough that it turns out to be of small importance for the model dynamics.

[20] Our definition of output excludes the value of imputed rents on housing services, which account for about 10% of GDP in the United States.

3.5. Housing and Renting

We assume that the minimum house size costs 1.5 (1.7) years of annual pre-tax (after-tax) household income: this constraint - together with the downpayment constraint - essentially prevents wealth-poor people from becoming homeowners too quickly.

Together with the minimum house size, the other parameter that has a large impact on the equilibrium homeownership rate is the utility penalty for renting. We calibrate this parameter in order to obtain a homeownership rate of 64 percent, the average value in the data for the period 1952-1982. The parameter that delivers this result is $\theta = 0.83$. This number implies that an agent is indifferent between renting a house of $1,000$ square feet and owning a house of 830 square feet (and would be happier if he owned a $1,000$ square feet house).

Our model has the property that, in absence of any utility penalty for renting, everyone would prefer renting, since renters do not incur transaction costs for changing house size, do not have to purchase a house of a minimum size, and do not have to save for the down-payment. (These outcomes are an equilibrium feature of model if we set $\theta = 1$).

4. Results

For expositional purposes, we break the presentation of the results in two parts. First, we illustrate the steady state properties of the model. Second, we illustrate the behavior and the time-series properties of the model economy in response to aggregate shocks.

4.1. Steady State Properties

4.1.1. General Features of Household Behavior

At each age, the household chooses consumption, saving, labor supply and housing investment taking account current (and expected) income, beginning of periods liquid assets and housing position. Here, we mostly focus on housing investment decisions, since other features of the model are in line with existing models of life-cycle consumption and saving behavior; we defer some illustration of the labor supply behavior to the next section, when we discuss the dynamics of the model in response to aggregate shocks.

It is simple to characterize the behavior of agents conditioning on whether they enter the period as renters or homeowners. For renters, the housing investment decision is relatively simple: at any given age, there is a threshold amount of liquid assets $\overline{a}$ ($-b$ in the notation of the model) such that, if $-b$ exceeds $\overline{a}$, households become homeowners. The larger initial liquid assets are, the less likely is a household to borrow against the house purchased.

For existing homeowners, there are four possibilities: homeowners can increase their house size, stay put, move down or switch to renting. Which option they choose depends on the combination of initial house and liquid assets they enter the period with. Figure 2 plots the optimal housing choice for a homeowner as a function of initial house size and liquid wealth.[21] The solid, downward sloping line plots the borrowing constraints that restricts debt not to exceed a fraction m of the housing stock. As the figure illustrates, larger liquid assets trigger larger investment in housing. In addition, purchasing and selling costs create a region of inaction where the household keep its housing stock constant. If liquid wealth falls, the household either moves to a smaller house or switches to renting.

An interesting feature of the model is that, for a household with very small liquid assets, the housing tenure decision is non-monotone 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 its house size, since, given the small amount of assets, the house size is closer to the optimal. If the initial house is medium-sized, the homeowner pays the adjustment cost and, because of the low liquid assets, switches to renting. If the initial house size is large, it is optimal to downsize, and to buy a smaller house.[22]

4.1.2. Life-Cycle Profiles

With these features of the model in mind, we show in Figure 3 the typical life-cycle of an agent in our model. We choose an agent with a low discount factor (and somewhat low assets) since the behavior of the agents close to the borrowing constraint is somewhat illustrative of the main workings of the model. This agent starts life with zero assets and rents the property where he lives. Over time, as his income rises because of life-cycle reasons, he consumes more and works less. Around age 38, he is hit by a negative income shock lasting one period only, that forces him to cut back on consumption and housing services. Hours fall by a very small amount. A sequence of good income shocks at age 40 and 41 allows him to save just enough to buy a house: in the period when he buys, the agent increases hours worked to save enough for the downpayment. After becoming a homeowner, hours move in opposite direction to wage shocks, rising in bad times, falling in good times (this mechanism is explained in detail in Section 4.2.1 below). Around age 60, the agent starts saving a bit for retirement. Upon retirement, the agent switches to a small rental property, and he dies at the age of 73.

An interesting dimension where it is useful to compare the model with the data is the frequency of housing adjustment for homeowners.[23] The work by Hansen (1998) based on the 1993 Survey of Income and Program Participation reports that the median homeowner stays in the same house for about 8 years. Anily, Hornik, and Israeli (1999) estimate that the average homeowner lives in the

[21] This figure is plotted for a patient agent at the age of 31 years, when individual and aggregate productivity and the average capital labor ratio are equal to their average value. A similar figure obtains for impatient agents.

[22] Smaller adjustment costs reduce the size of the inaction region. In the limiting case of no adjustment costs, the inaction region disappears.

[23] In the model, because there are no costs associated with modifying housing consumption, renters change their housing position every period. This assumption is in line with the data, that show that on average renters move about every 2 years.

same residence for 13 years. The corresponding number for our model (when calibrated to match the 1990s) is 12 years. In the data, 15 (35) percent of the moves are associated with a move to a different state (county), so they are probably associated with "moving shocks" that essentially force owners to sell and buy a house even at constant housing consumption. Excluding these moves, one can realistically conclude that in the data homeowners change their housing consumption on average every 10 years, a number that is similar to the model.

Figure 4 compares the age profiles generated by the model with those coming from the data. (Figure 5 reports the analogous profiles for the calibration meant to capture the 1990s).

4.1.3. The Wealth Distribution

Our model replicates the U.S. wealth distribution almost exactly. The Lorenz curves for the U.S. economy and for our model economy are reported in Figure 6. The Gini coefficient for wealth in the model is 0.75, and is virtually the same as in the data (equal to 0.78, according to Budria Rodriguez et al., 1997). The main discrepancy between the model and the data is that we underestimate the fraction of wealth held by the top 5% of the population. On the positive side, the model does remarkably well at matching the fraction of wealth held by the poorest 40 percent of the U.S. population, which has essentially no assets and no debt, like the renters of our model.

In this dimension, a model without preference heterogeneity (see Section 6) would do much worse. The Gini coefficient for wealth in the model with a single discount factor is 0.52, much lower than in the data.[24]

4.2. Business Cycle Results

In this section, we try to shed light on the propagation mechanism of aggregate shocks. Unlike the standard representative agent (real business cycle) model, heterogeneity in this context will imply that individuals will respond differently to common shocks. Here, there are two aspects of heterogeneity that matter: one is purely exogenous, and reflects the assumption that individuals with different ages have different productivity, planning horizon, and utility weights: because other papers[25] have extensively studied these features in life-cycle models with aggregate shocks, we do not devote much space to these issues here. Instead, we focus on the endogenous component of heterogeneity, which reflects the fact that individuals with different ages and income histories accumulate over time different amounts of wealth (housing and non-housing assets): in turn, heterogeneity in wealth implies different individual responses to the same shock. It is at these responses that we turn now in order to illustrate the workings of our model.

[24] We recalibrate the model with a single discount factor to obtain the same homeownership rate as in our baseline economy.

[25] See for instance the work of Rios-Rull (1996) and Gomme et al. (2004).

4.2.1. Workings of the Model

To describe the workings of the model, we focus on the response of aggregate hours to a positive shock to productivity. In general equilibrium, the shock changes the incentives to work and to save of all agents. Here we take a partial equilibrium approach to show how the wealth distribution and its composition affect the responses of the agents in our model. To do so, we consider a stripped-down version of the budget constraint of an individual that keeps total beginning of period wealth ξ constant between two periods. Abstracting from pensions, this implies the following budget constraint:

$$c_t = \underbrace{w_t \eta_a z_t l_t}_{\text{wage income}} + \underbrace{\xi}_{\text{interest income net of housing maintenance}}$$

where $\xi = -(R-1)\,b - \delta_h h$ measures the resources besides wage income that can be used for consumption.[26] The term $-(R-1)\,b$ denotes net interest income, the term $\delta_h h$ is the maintenance cost that is required to keep the housing value unchanged.

Given this constraint, different values of ξ can be mapped into different positions of the agents along the wealth distribution. For a wealthy homeowner, ξ is positive and large, and wage income is a small fraction of non-durable consumption c. For a renter, h equals 0 by definition; 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 constant ξ, the log-linearized budget constraint becomes, denoting with $\widehat{x} \equiv \frac{x_t - x}{x}$, where x is the steady state value of a variable:

$$\widehat{c}_t = \frac{wl}{c}\left(\widehat{w}_t + \widehat{l}_t\right) \text{ or } \widehat{l}_t = \widehat{w}_t + \frac{c}{wl}\widehat{c}_t.$$

This version of the constraint can be interpreted, given the wage, as an equation dictating how much the household needs to work to finance a given consumption stream (we call this equation the "labor need" curve). The larger the desired change in consumption $\widehat{c}$, the larger the required hours $\widehat{l}$ needed to finance consumption, with an elasticity of hours to consumption that is given by the ratio of consumption to wage income $\frac{c}{wl} \equiv \phi$. For a wealthy homeowner, ϕ is high and larger than unity, since labor income is a small share of its total income stream; for a renter without assets, $\phi = 1$; finally, for an indebted homeowner, $\phi < 1$, reflecting the need to use part of the income to finance maintenance cost and to pay back the mortgage. In words, a wealthy guy needs to increase hours worked by more than 1% to finance a 1% rise in consumption, since labor income makes for less than 100% of its consumption; an indebted homeowner needs to increase hours worked by less than 1% to finance a 1% rise in consumption, because of the leverage effect; a renter needs to increase

[26] For a renter, housing and non-housing expenditures are proportional to each other, so the constraint reads

$$c_t = \frac{1}{1+j}\left(w_t \eta_a z_t l_t + \xi\right)$$

where j is the optimal ratio of housing expenditure to non-durable consumption. With minor modifications, the arguments in this section carry over to this case, since ξ cannot be negative for renters.

hours 1 for 1 with consumption. These relationships are plotted in Figure 7 for the three types. The labor need curve is upward sloping, with a slope given by the consumption/wage ratio ϕ. In the consumption-hours space, this curve is shifted down and to the right by a rise in the wage, since a higher wage increases the consumption possibilities for given hours worked.

The other key equation determining hours is the traditional labor supply curve. After manipulation, this curve reads as:

$$\widehat{l} = \varepsilon \left(\widehat{w} - \widehat{c} \right)$$

where ε is the steady state Frisch labor supply elasticity. This curve slopes downward because of the wealth effect on labor, and is shifted to the right by a rise in the wage.

Figure 7 also shows how the labor supply and labor need curves move for a given change in the wage, say 1 percent. For illustrative purposes, we take the change in the wage as the exogenous driving force of the model here, since an exogenous rise in productivity exerts a direct effect on the wage. The rise in the wage shifts the labor need curve down by 1% and the labor supply curve up by ε%. If ε is larger than one, as in our baseline calibration,[27] borrowers (savers) are more likely to reduce (increase) hours following a positive wage shock, because their wealth effect more than offsets (does not offset) the labor need effect.

In the aggregate, consider aggregate hours broken down as follows:

$$\underset{\text{total hours}}{L_t} = \underset{\text{hours of renters}}{\int_{h=0} l_a \left(x_t \right) \eta_a z_t \partial \Phi_t} + \underset{\text{hours of indebted homeowners}}{\int_{h>0, b>0} l_a \left(x_t \right) \eta_a z_t \partial \Phi_t} + \underset{\text{hours of positive-wealth homeowners}}{\int_{h>0, b\leq 0} l_a \left(x_t \right) \eta_a z_t \partial \Phi_t} .$$

Denote n_R, n_{OD} and n_{OC} the fraction of renters, indebted homeowners and positive-wealth homeowners respectively. The average percentage response of hours to a wage shock can be approximated as a weighted average of the responses of the three types:

$$\widehat{L} = n_R \underset{\text{zero}}{\widehat{L}_R} + n_{OD} \underset{\text{negative}}{\widehat{L}_{OD}} + n_{OC} \underset{\text{positive}}{\widehat{L}_{OC}} ,$$

where $\widehat{L}_i$ is the average hours worked by household category i, $i = R, OD, OC$. For the economy as a whole, the aggregate response of hours to a productivity shock is positive: this happens because total capital income is positive, so that substitution and wealth effects do not offset with each other. However, as we described above, there is wide heterogeneity in the hours worked across segments of the population: structural changes in the economy that affects the distribution of wealth can produce aggregate business cycle effects.

4.2.2. Business Cycle Statistics of the Model

We begin this section with a brief recap of the empirical regularities concerning housing investment, debt and economic activity that are most relevant to our analysis. In quarterly, HP-filtered, postwar

[27] Traditional RBC models require or assume labor supply elasticities which are greater than unity. See King and Rebelo (1999).

US data, the relative variability of housing investment is large, with a standard deviation that is between 3 and 4 times that of GDP. Next, housing investment is procyclical, with a correlation with GDP around 0.8 (in the period 1952-1982). Taken together, these two facts imply that the growth contribution of housing investment to the business cycle is larger than its share of GDP. Another important aspect of the data is that household mortgage debt is strongly procyclical in data from 1952 to 1982, although it becomes essentially acyclical thereafter, with a correlation with GDP that drops from 0.78 to 0.16.

Against this data background, Table 2 reports some of the key statistics generated by the benchmark model, and compares them to the data. Overall, our baseline model appears to do a good job in reproducing the relative volatility of each component of aggregate demand, including housing investment. In particular, the model can account for about 60 of volatility of housing investment. On the contrary, the model tends to overpredict the volatility of aggregate consumption. The volatility of business investment is smaller than in data, but, as in the data, it is smaller than the volatility of housing investment: this happens despite the fact that our model assumes that adjusting housing capital is costly, whereas there are no costs for changing business capital.

Turning to household debt, the model does quite well in reproducing the cyclical behavior of household debt (although it overpredicts its volatility). Key to this result is the fact that a large bulk of the population (the impatients) upgrades its housing in good times taking a mortgage at the same time, in particular the very young and the middle aged.

5. Dynamic Consequences of Higher Idiosyncratic Volatility and Lower Downpayments

Having shown above that the model roughly captures postwar US business cycles, we now consider the implications of two experiments. In the first, we lower the down-payment requirement from 25 to 15 percent. In the second, we increase the amount of idiosyncratic risk faced by households. When both changes are active, our experiment is intended to mirror two of the main structural changes that have occurred in the U.S. economy since the Great Moderation of the mid 1980's. The results are shown in Table 3.

5.1. A Decrease in Downpayment Requirements

A larger value of m_h has two main effects on the properties of the model: (1) it leads to an increase in the homeownership rate; and (2) reduces the volatility of household investment and, to a lesser extent, of the other components of demand.

Lower downpayments allow more housing ownership among those with very little net worth. This is illustrated in Figure 8, that depicts the average life-cycle profiles of homeownership and housing wealth separately for booms and recessions, with a high down-payment (blue/solid lines) and a low down-payment (red/dashed line). The prediction of the model is that homeownership increases substantially for those who are between the ages of 30 and 65.

Turning to business cycles, the rise in m_h reduces the volatility of housing investment. Why? There are three main forces at work. One works directly though a housing demand-side channel. When downpayments are high, more agents are unable to save enough for the downpayment, or save just enough to afford the minimum house size. The housing investment of these agents strongly reacts to shocks: they switch from renting (smaller house) to owning (bigger house) in good times, and from owning to renting in bad times. With a lower down-payment (higher m_h) instead, less people are constrained and adjustments in their housing stock are in general smoother over the life-cycle, occurring independently of the business cycle fluctuations. As a result, a lower down-payment leads to lower volatility of housing investment.

The second force that reduces the volatility of housing investment mainly operates through the supply side. As we explained above, indebted homeowners (compared to renters) are more likely to reduce hours worked in response to technology shocks, so their presence acts as a dampener of aggregate technology shocks. Therefore, the higher homeownership rate induced by looser borrowing constraints contributes to lower aggregate volatility.

The third force is somewhat intuitive. On average, because of adjustment costs, homeowners modify their housing consumption very little over time, relative to renters. This reduces the volatility of housing investment more than the other components of expenditure.

5.2. An Increase in Earnings Volatility

An increase in earnings volatility slightly reduces homeownership. The lower homeownership rate would tend to increase the volatility of total housing investment, leaving unaffected the other properties of the model. However, when changes in earnings volatility are coupled together with lower downpayments, higher volatility tends to reinforce the effects of lower downpayments on the volatility of aggregates.

5.3. Combining Lower Downpayments and Higher Volatility

When lower downpayments (higher m_h) and higher volatility (higher σ_Z) are combined together, they replicate the observed increase in homeownership rates observed in the United States over the last 40 years: home ownership rates in the US were around 64 percent in the 1960 and the 1970s, and rose in the 1990s to an average of 67 percent.

The last column of Table 3 also shows the business cycle consequences of these two structural changes: here, two interesting results emerge. First, the volatility of housing investment (and, to a lesser extent, of GDP) falls by more than would be predicted by changing the two parameters in isolation. Second, the combined effect of these two forces makes aggregate debt acyclical.

Our interpretation for these results is as follows: in response to a combination of high leverage (induced by lower downpayments) and higher income volatility, leveraged individuals become more *cautious* in response to aggregate shocks, thus reducing the extent to which they change their borrowing and their housing demand in response to aggregate shocks. (This is especially true for housing since housing purchases involve durable goods and are subject to adjustment costs.) Using Figure 2 as the reference point, the area of inaction where a household does not change its housing stock in response to changes in wealth becomes larger (at the expense of the "renting" and the "moving down" regions), especially for very low levels of liquid assets. This result reflects two forces: higher volatility reduces the willingness to change housing consumption too much in presence of adjustment costs; lower downpayments allow households not to switch to smaller houses or to renting in bad lines, since they provide a larger buffer against bad income shocks.

Taken together, higher volatility and lower downpayments change the cyclical behavior of household debt. The model correlation between debt and GDP falls from 0.88 to 0.10. They also lead to a noticeable decrease in the volatility of housing investment: its standard deviation falls from 4.7 to 4.1 percent.

6. Robustness Analysis

No Housing Adjustment Costs. A model without housing adjustment costs delivers a volatility of housing investment which is slightly higher than in the baseline model. It is 2.24 percent in the baseline case, 2.34 without housing adjustment costs.

Everybody rents. A model in which everybody rents ($\theta = 1$) behaves very similarly to a model without housing adjustment costs. The main difference is that in this model there is no debt in equilibrium. In this model, the standard deviation of residential investment is 2.88, whereas it is 2.15 in the baseline model.

Everybody owns. If we set the penalty for renting to its maximum value, so that $\theta = 0$, everybody saves enough to accumulate for the downpayment and the equilibrium homeownership rate is 100 percent. Because of the large individual adjustment costs associated with changing housing positions, the volatility of housing investment is now smaller, at 1.43 percent.

Homogeneous Discount Factor. To show the properties of the model with a homogeneous discount factor, we recalibrate our model economy so that it has the same homeownership rate and interest rate as in our baseline model. We change two parameters, the discount factor (previously 0.925 for 2/3 of the population, and 0.995 for 1/3) and the relative utility from renting (previously at 0.83). The values that achieve same interest rate and homeownership rate are $\beta = 0.98$ and $\theta = 0.95$. At these parameter values, the volatility of housing investment is higher than in the baseline calibration, at 2.53 percent. The volatility of total output is also much higher, 2.24 versus 2.09 percent. It is not difficult to see why. With a single discount factor, very few people hold debt in equilibrium, and the distribution of wealth is much more egalitarian than in the data. Because of this, the average elasticity of labor supply is higher than in an economy with lots of debtors. However, this model fails to account for the inequality of wealth in data. The Gini coefficient for wealth is in fact 0.52, far lower than in the data.

7. Conclusions

In this paper, we generalize an equilibrium business cycle model to allow for endogenous housing investment. We model a house as a big, lumpy item that can be purchased or rented and that can be adjusted at a cost. The resulting dynamics of housing investment are realistic not only at the macroeconomic level, but also at the level of individual household behavior: even if individuals adjust housing only infrequently, housing investment turns out to be the most volatile component of aggregate demand, in our stylized model as in the data.

We show that our model accounts for the procyclicality of housing investment and for a good part its relative variability. At the same time, our model can explain the procyclical behavior of household debt. Even in absence of the microfoundations of housing demand that our model embeds, these results have often eluded existing macroeconomic models of housing investment and portfolio choice.

References

[1] Anily, Shoshana, Jacob Hornik, and Miron Israeli (1999). "Inferring the Distribution of Households' Duration of Residence from Data on Current Residence Time." Journal of Business and Economic Statistics 17(3): 373–81.

[2] Budria Rodriguez S., Diaz-Gimenez J., Quadrini V. and Rios-Rull J.V. (2002), "Updated facts on the U.S. distributions of earnings, income, and wealth," Quarterly Review, Federal Reserve Bank of Minneapolis, issue Sum, pp. 2-35.

[3] Cagetti M. (2003), "Wealth Accumulation over the Life Cycle and Precautionary Savings," *Journal of Business and Economic Statistics*, vol. 21(3), pp. 339-53.

[4] Campbell J. R. and Hercowitz Z. (2006), "The Role of Collateralized Household Debt in Macroeconomic Stabilization," Working Paper

[5] Carroll C. (2000), "Requiem for the Representative Consumer? Aggregate implications of Microeconomic Consumption Behavior," *American Economic Review Papers and Proceedings*, vol. 90 (2), pp. 110-115

[6] Davis, Morris A., and Heathcote, Jonathan. "Housing and the Business Cycle." *International Economic Review,* August 2005, *46*(3), pp. 751-784.

[7] Dynan K., Elmendorf D.W. and Sichel D. (2007), "The Evolution of Household Income Volatility," Finance and Economics Discussion Series n.61, Board of Governors of the Federal Reserve System

[8] Fernandez-Villaverde, Jesus and Dirk Krueger (2005), "Consumption and Saving over the Life Cycle: How Important Are Consumer Durables?", Working Paper, University of Pennsylvania.

[9] Fisher, Jonas D. M. (2007) "Why Does Household Investment Lead Business Investment over the Business Cycle?." *Journal of Political Economy*, *115*(1), pp. 141-68.

[10] Fisher Jonas D. M. and Gervais M. (2007), "First-Time Home Buyers and Residential Investment Volatility," Working Paper no. 2007-15, Federal Reserve Bank of Chicago

[11] Gervais M. (2002), "Housing Taxation and Capital Accumulation,", *Journal of Monetary Economics*, vol. 49, pp. 1461-1489.

[12] Gottschalk P. and Moffitt R.A. (2002), "Trends in the Transitory Variance of Earnings in the United States," Economic Journal, vol. 412 (478), pp. C68-C73.

[13] Gomme P., Rogerson R., Rupert P. and Wright R. (2004), "The Business Cycle and the Life Cycle," NBER Macroannual, pp. 415-461.

[14] Gruber, Joseph W., and Robert F. Martin (2006). “Precautionary Savings and the Wealth Distribution with Illiquid Durables.” International Finance Discussion Paper 773.

[15] Haider S.J. (2001), "Earnings Instability and Earnings Inequality of Males in the United States, 1967-1991," *Journal of Labor Economics,* vol. 19 (4), pp. 799-836.

[16] Hansen, Kristin A. (1988), "Seasonality of Moves and Duration of Residence," p70-66, U.S. Census Bureau, October.

[17] Hansen, G. D. (1993), "The Cyclical and Secular Behaviour of the Labour Input: Comparing Efficiency Units and Hours Worked," *Journal of Applied Econometrics*, vol. 8(1), pp. 71-80.

[18] Heathcote J., Storesletten K. and Violante G.L. (2008), "The Macroeconomic Implications of Rising Wage Inequality in the U.S.", mimeo, NYU

[19] Heaton, J., and D. J. Lucas (1996), "Evaluating the Effects of Incomplete Markets on Risk Sharing and Asset Pricing.", *Journal of Political Economy*, 104(3), pp. 443-87.

[20] Hendricks L. (2007), "How Important Is Discount Rate Heterogeneity for Wealth Inequality?," *Journal of Economic Dynamics and Control,* vol. 31 (9), pp. 3042-3068.

[21] King, Robert G., and Sergio T. Rebelo (1999), Resuscitating real business cycles. In: John B. Taylor and Michael Woodford, Editor(s), Handbook of Macroeconomics, Elsevier, Volume 1, Part 2, Pages 927-1007.

[22] Kiyotaki, Nobuhiro, Alexander Michaelides and Kalin Nikolov (2007), “Winners and Losers in Housing Markets”, mimeo.

[23] Krueger, Dirk, and Fabrizio Perri (2006), “Does Income Inequality Lead to Consumption Inequality? Evidence and Theory,” *Review of Economic Studies*, 73, 1, 163-194.

[24] Krusell P. and Smith A.A.J. (1998), "Income and Wealth Heterogeneity in the Macroeconomy", *Journal of Political Economy*, vol. 106 (5), pp. 867-896.

[25] Ljungqvist, L. and T.Sargent (2004), Recursive Macroeconomic Theory, 2nd edition, The MIT Press, Cambridge, MA.

[26] Nakajima, Makoto. (2005), ”Rising Earning Instability, Portfolio Choice and House Prices”, Working Paper, University of Illinois, Urbana-Champaign.

[27] Ortalo-Magne, Francois, and Sven Rady (2006). “Housing Market Dynamics: On the Contribution of Income Shocks and Credit Constraints,” Review of Economic Studies, vol. 73(2), pages 459-485, 04.

[28] Poterba J.M. (1992), "Taxation and Housing Markets," in Shoven J. and Whalley J. (eds.), *Canada-U.S. Tax Comparisons*, Chicago: University of Chicago Press, pp. 275-295.

[29] Rios-Rull, Jose-Victor, (1996). "Life-Cycle Economies and Aggregate Fluctuations," Review of Economic Studies, Blackwell Publishing, vol. 63(3), pages 465-89, July.

[30] Rosen H.S., "Housing Subsidies: Effects on Housing Decisions, Efficiency, and Equity," in Alan J. Auerbach and Martin Feldstein, eds., Handbook of Public Economics, vol. 1. Amsterdam, North Holland, 1985, pp. 375-420.

[31] Scholz, John Karl, Ananth Seshadri, and Surachai Khitatrakun. 2006. "Are Americans Saving 'Optimally' for Retirement?" *Journal of Political Economy*, 114(4), August: 607-643.

[32] Silos P. (2005), "Housing Tenure and Wealth Distribution in Life-Cycle Economies," Working Paper no. 2005-25, Federal Reserve Bank of Atlanta

[33] Silos P. (2007), "Housing, Portfolio Choice, and the Macroeconomy," *Journal of Economic Dynamics and Control*, vol. 31 (8), pp. 2774-2801.

[34] Storesletten, K., Telmer, C.I., Yaron, A., (2004). Cyclical dynamics in idiosyncratic labor market risk. Journal of Political Economy 112 (3), 695–717.

[35] Tauchen G. (1986), "Finite State Markov-chain Approximations to Univariate and Vector Autoregressions," *Economics Letters*, vol. 20, pp. 177-181.

[36] Venti S.F. and Wise D.A. (2001), "Choice, Chance, and Wealth Dispersion at Retirement," in *Aging Issues in the United States and Japan*, Ogura S., Tachibanaki T. and Wise D.A. (eds.), Chicago, University of Chicago Press, pp. 25-64.

Appendix A: Computational Details

We numerically solve for the model equilibrium using a computational method similar to the one used in Krusell and Smith (1998). The value and policy functions are computed on grids of points for the state variables, and then approximated with linear interpolation at points not on the grids. The algorithm consists of the following steps:

1. Specify grids for the state space of individual and aggregate state variables.

   As a result of robustness checks, the number of grid points was chosen as follows: 5 points for the aggregate shock, 3 values for the idiosyncratic shock, 16 points for the housing stock, and 300 points for the financial asset.[28] For aggregate capital to labor ratio, we initially choose a grid of 18 equally spaced points in the range $[0.8(K/L)^*, 1.2(K/L)^*]$, where $(K/L)^*$ denotes the average value of this variable in the simulations. The grid is then updated at each iteration consistently with the simulated K/L, assigning as its boundaries the minimum and the maximum simulated values.

2. Guess initial coefficients $\left\{\omega_i^A\right\}_{A\in\widetilde{A},i=0,1}$ for the linear function that approximates the law of motion of the aggregate capital-labor ratio:

$$K_t/L_{t+1} = \omega_0^A + \omega_1^A K_{t-1}/L_t$$

3. Starting from age T backward, compute optimal policies as a function of the individual and aggregate states, solving first the homeowner's and renter's problems separately.[29] Notice that the intra-temporal optimal value for labor hours as a function of consumption and productivity shock for ages $a \leq \widetilde{T}$ is the following:[30]

$$l_{a,t} = \bar{l} - \frac{\tau c_{a,t}}{w_t \eta_a z_t}$$

   which allows one to derive consumption before age $\widetilde{T}$ directly from the budget constraint. For the homeowner:

$$c_{a,t} = \frac{w_t \eta_a z_t \bar{l} - R_t b_{a,t-1} + b_{a,t} + (1-\delta_H) h_{a,t-1} - h_{a,t} - \Psi(h_{a,t}, h_{a,t-1})}{1+\tau}$$

   so that the per-period utility function for $a \leq \widetilde{T}$ can be transformed as follows:

$$\widetilde{u}(c_{a,t}, h_{a,t}, w_t z_t) = (1+\tau)\log c_{a,t} + j \log h_{a,t} + \tau \log(\tau / w_t \eta_a z_t)$$

   For the tenant, taking into consideration the intra-temporal condition for optimal house services to rent:

$$c_{a,t} = \frac{w_t \eta_a z_t \bar{l} - R_t b_{a,t-1} + b_{a,t} + (1-\delta_H) h_{a,t-1} - \Psi(0, h_{a,t-1})}{1+\tau+j}$$

   so that the per-period utility function for $a \leq \widetilde{T}$ can be transformed as follows:

$$\widetilde{u}(c_{a,t}, p_t^s, w_t z_t) = (1+\tau+j)\log c_{a,t} + j \log(j\theta / p_t^s) + \tau \log(\tau / w_t \eta_a z_t)$$

[28] The upper bound for the housing grid and the lower bound for debt are chosen to be wide enough so that they never bind in the simulations.

[29] In computation, we exploit the strict concavity of the value function in the choice for assets as well as the monotonicity of the policy function in assets (for the homeowner problem, the monotonocity is for any given choice of the housing stock).

[30] We prevent individuals from choosing negative hours.

As a consequence, the homeowner's dynamic optimization problem entails solving for policy functions for b and h only, while the renter's one consists in solving for b only. The problems of the retired people $(a > \widetilde{T})$ are similar to the above, where we set $\tau = 0$.

4. Draw a series of aggregate and idiosyncratic shocks according to the related stochastic processes. Draw a series of "death" shocks according to the survival probabilities. Use the (approximated) policy functions and the predicted aggregate variables to simulate the optimal decisions of a large number of agents for many periods (We solve for the value and policy functions on a finite number of grid points for b, h and K/L. In the simulations, we perform linear interpolation between grid points for b', but we restrict the choices of h' to lie on the grid.) We simulate 50,000 individuals for 5,000 periods, discarding the first 500.[31] Compute the aggregate variables K_{t+1}/L_t at each t.

5. Run a regression of the simulated aggregate capital-labor ratio on past aggregate capital-labor ratio, retrieving the new coefficients $\{\omega_i^A\}$ for the law of motion for K_{t+1}/L_t. We repeat steps 3 and 4 until convergence over the coefficients of the regressions. We measure goodness of fit using the $R - squared$ of the regressions (which are always equal to 0.99 or higher at convergence).

[31] We enforced the law of large numbers by making sure that the simulated fractions of ages and of labor productivity shocks correspond to the theoretical ones, by randomly adjusting the values of the shocks.

Appendix B: Calibrating the income process

The level of the volatility

The (parsimonious) process for individual income productivity that we specify in the model is:

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

We want to have values for ρ_Z and σ_Z that are in line with micro evidence. Below is a survey of studies that have attempted to estimate these parameters in setups similar to ours.[32]

1. Heaton and Lucas (1996) consider a model where *family* log labor income y_t (normalized by total labor income) evolves according to a first-order autoregression of the form similar to ours. We report their findings using our notation. After controlling for fixed effects (η), for the period 1969 to 1984 (see their Table A.3), they estimate $\rho_Z = 0.53$ and $\sigma_Z = 0.296$.

$$y_t = \eta + 0.53 y_{t-1} + 0.296 \left(1 - 0.53^2\right)^{1/2} \varepsilon_t$$

2. Scholz, Seshadri, and Khitatrakun (2006) specify and estimate a model of household log labor earnings that controls for fixed effects, a polynomial in age, and autocorrelation in earnings. Their sample is the social security earnings records (see Section 1.B). Their estimates (see their Table B of Appendix A) for married, no college, two-earners are as follows:

$$y_t = \eta + g\left(age\right) + 0.699 y_t + 0.428 \left(1 - 0.7^2\right)^{1/2} \varepsilon_t$$

3. Storesletten, Telmer and Yaron (2004) use PSID data from 1968 to 1993 to estimate household-level income process with persistent and transitory income shocks, so that

$$y_t = \rho y_{t-1} + \varepsilon_t^T + \varepsilon_t^P - \rho \varepsilon_{t-1}^P$$

and ε_t^T is an iid shock, ε_t^P is persistent shock. Their estimated (see Table 2 in their paper) standard deviation for ε_t^P, the standard deviation of the persistent component, is 0.162 in expansions, 0.088 in recessions (0.125 on average). Hence this gives, setting aside the purely transitory shocks:

$$y_t = 0.963 y_{t-1} + 0.463 \left(1 - 0.963^2\right)^{1/2} \varepsilon_t^P \text{ (Exactly identified model, A)}$$

In the Overidentified, CRSP Model (Model F), they have 0.159 and 0.084 (0.121 on average). This gives

$$y_t = 0.939 y_{t-1} + 0.351 \left(1 - 0.939^2\right)^{1/2} \varepsilon_t^P \text{ (model F).}$$

The Change in Volatility

1. Using PSID data, Heathcote, Storesletten, and Violante (2008) decompose the evolution of the cross-sectional variance of individual earnings over the period 1967-1996 into the variances of fixed effects, persistent shocks, and transitory shocks. In particular, they find that the variance of persistent shocks doubles during the 1975-1985 decade.

[32] One difference between our productivity process and the estimates of wage processes from the literature is that the two only coincide if hours are constant over the life cycle.

2. Gottschalk and Moffitt (2002) perform a similar decomposition: they find that the increase in the variance of the non-permanent component of log earnings is 10 basis points, from 0.31 to 0.41.

3. Haider (2001) finds that increases in earnings instability over the 1970s and increases in lifetime earnings inequality in the 1980s account in equal parts for the increase of inequality in the data. To measure the magnitude of earnings instability in year t, he uses the cross-sectional variance of the idiosyncratic deviations in year t. His estimate of ρ is 0.639 (Table 4). He finds that the unconditional standard deviation of the instability component goes from around $0.23 - 0.24$ to about $0.35 - 0.37$.

4. Krueger and Perri (2006) model income as an ARMA process. In 1980, they find:

$$\begin{aligned} y_t &= z_t + \varepsilon_t \\ z_t &= 0.9989 z_{t-1} + 0.42 \left(1 - 0.9989^2\right)^{1/2} \varepsilon_t^z \\ \varepsilon_t &= 0.28 \varepsilon_t^e \end{aligned}$$

   In 2003, they have:

$$\begin{aligned} z_t &= 0.9989 z_{t-1} + 0.52 \left(1 - 0.9989^2\right)^{1/2} \varepsilon_t^z \\ \varepsilon_t &= 0.36 \varepsilon_t^e \end{aligned}$$

   The increase in total variance of income is 20 basis points, from 0.33 to 0.53. That corresponds to a 15 basis points increase in the standard deviation of log income.

5. Dynan, Elmendorf and Sichel (2007) survey studies and use independent evidence to reach similar conclusions. In particular, they estimate that the standard deviation of percent changes in household income rose one-fourth between the early 1970s and the early 2000s. They model household earnings as a random walk, so their method is not directly comparable with ours. In any event, they report that the standard deviation of the innovations to income growth rose from 0.2 to 0.25, approximately.

Summary

From this brief survey, we conclude that a plausible value for the persistence of the income shock is 0.9. We set the standard deviation of income to be equal to 0.3 in the early part of the sample, which is the lower bound of the estimates reported above. We set the standard deviation to 0.45 in the second part of the sample: a change of 0.15 is the upper bound of the estimates for the increase in the volatility.

Appendix C: Notes on Aggregate Volatility

Conventional studies of the Great Moderation use as a reference variable chain-weighted GDP and its chain-weighted components. This way of treating the data is perhaps the most appropriate when one cares about volatility of real quantities.

An alternative way of looking at the moderation is to consider each component of demand scaled by the GDP deflator. That is, consider this definition of nominal GDP:

$$PY = P_C C + P_I I$$

fixed-weight GDP is:

$$Y = \frac{P_C}{P} C + \frac{P_I}{P} I$$

In the data, the volatility of $\frac{P_I}{P} I$ has fallen far less than the volatility of I. Evidently, one explanation is that the volatility of $\frac{P_I}{I}$ has risen, or that the covariance of $\frac{P_I}{P}$ with I has gone up, so that omitting $\frac{P_I}{I}$ leads to overestimate the decline in the volatility of residential investment.

Table 1. US Economy. Cyclical Statistics.

	1952.I -1982.IV (Early Period)			1984.I -2008.II (Late Period)		
	stdev %	ratio[i]	corr. w/ GDP	stdev%	ratio[i]	corr. w/ GDP
GDP[ii]	2.09	1.00	1.00	1.20	1.00	1.00
C[iii]	1.20	0.57	0.92	0.58	0.48	0.91
IH	8.24	3.94	0.84	3.61	3.01	0.66
IK	5.03	2.41	0.77	4.08	3.41	0.85
$Debt$	2.23	1.21	0.78	1.56	1.31	0.16

Notes: (*i*) The ratio is the standard deviation of the variable divided by that of GDP; (*ii*) C, IH and IK are chain-weighted consumption, residential investment and business investment respectively; GDP is the sum of the nominal series divided by the GDP deflator; (*iii*) Consumption of durables is assigned to IH, not to C. All series are in logs and detrended with HP-filter with smoothing parameter 1,600. Prior to detrending, the series are scaled by civilian non-institutional population.

Table 2: US Economy and Model. Cyclical Statistics. Comparison for the Early Period.

	DATA: 1952.I -1982.IV			MODEL		
	stdev %	ratio[i]	corr. w/ GDP	stdev%	ratio[i]	corr. w/ GDP
GDP[ii]	2.09	1.00	1.00	2.09	1.00	1.00
C[iii]	1.20	0.57	0.92	1.69	0.81	0.95
IH	8.24	3.94	0.84	4.68	2.24	0.90
IK	5.03	2.41	0.77	3.50	1.68	0.86
$Debt$[iv]	2.23	1.06	0.78	7.77	3.72	0.88

Notes: (*i*) The ratio is the standard deviation of the variable scaled by that of GDP; (*ii*) C, IH and IK are chain-weighted consumption, residential investment and business investment respectively; GDP is the sum of the nominal series $(C + IH + IK)$ divided by the GDP deflator; (*iii*) Durables are considered part of IH, not of C; (*iv*) Debt is gross household mortgage debt outstanding deflated by the GDP deflator. Prior to detrending (with HP-filter), all series are scaled by civilian non-institutional population. The model series are based on simulations of 5,000 periods and are HP-filtered. The weight for the HP-filter is 1,600 for the quarterly data, 6.25 for the model (each model period is a year).

Table 3: Model predictions, changing downpayment requirements and income volatility

	Baseline (1)	(2)	(3)	(4)
	Early Period			Late Period
	$m = 0.75$	$m = 0.85$	$m = 0.75$	$m = 0.85$
	$\sigma_Z = 0.3$	$\sigma_Z = 0.3$	$\sigma_z = 0.45$	$\sigma_Z = 0.45$
stdev%				
GDP[ii]	2.092	2.040	2.093	2.038
C[iii]	1.69	1.67	1.70	1.70
IH	4.68	4.39	4.78	4.12
IK	3.50	3.24	3.53	3.34
$Debt$[iv]	7.77	2.22	6.72	1.72
Corr$(Debt, GDP)$	0.88	0.52	0.80	0.10
% $Homeown$	64%	75%	63%	67%
Debt to GDP	0.31	0.46	0.27	0.35
$\sigma\,(\Delta C_{it}/C_i)$	0.14	0.14	0.17	0.17
Gini wealth	0.76	0.76	0.76	0.77
Gini labor income	0.43	0.43	0.49	0.49
Gini consumption	0.26	0.26	0.30	0.31

Baseline calibration and sensitivity analysis. Model 1 is the baseline model that is targeted to the U.S. data for the period 1952-1982. Model 2 increases the loan-to-value ratio from 0.75 to 0.85. Model 3 increases earnings volatility from 0.3 to 0.45. Model 4 increases both loan-to-values and earnings volatility and targets the U.S. economy for the period 1984-2008.

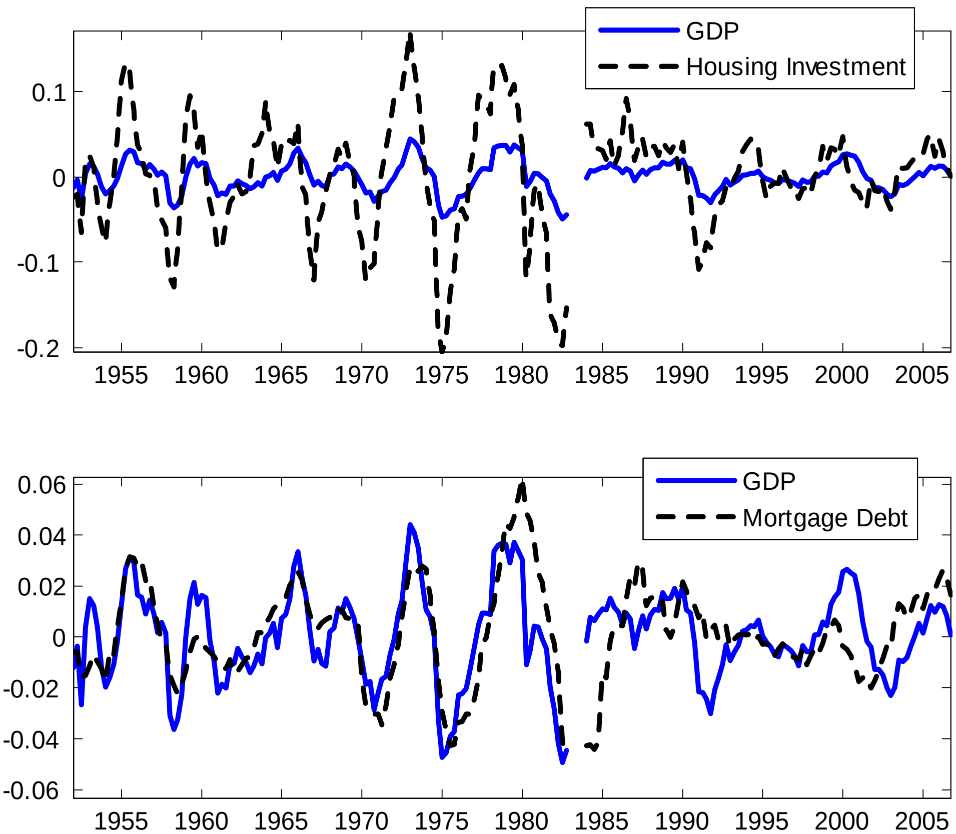

Figure 1: Household debt, Housing Investment and GDP (HP-filtered variables).

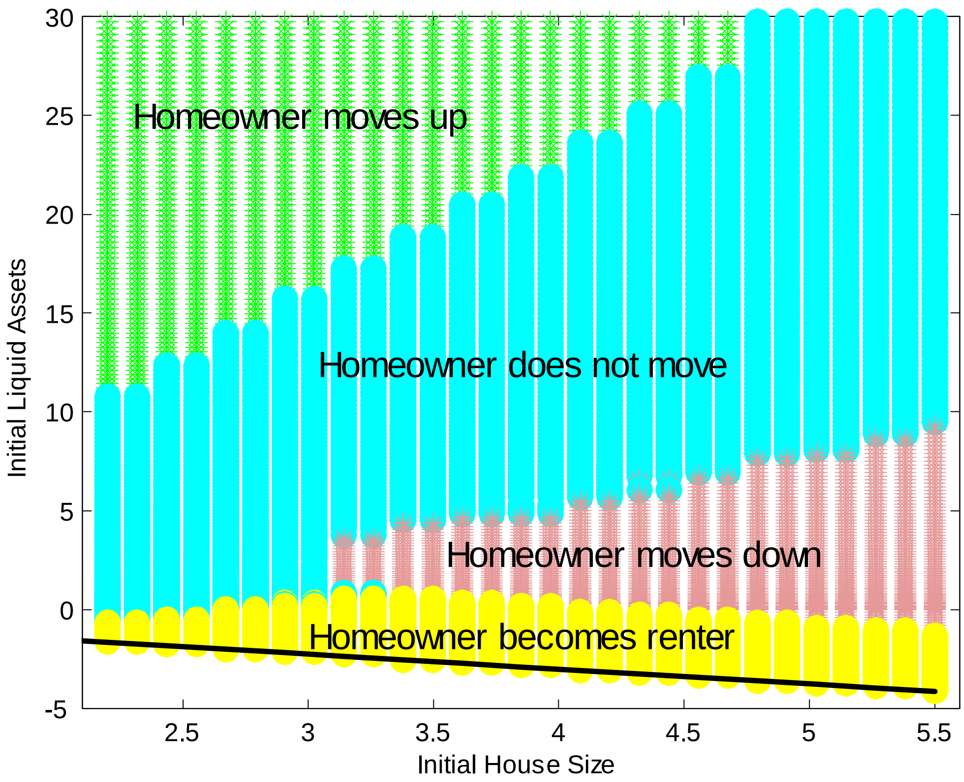

Figure 2: Homeowner's housing investment decision as a function of initial house size and liquid assets.

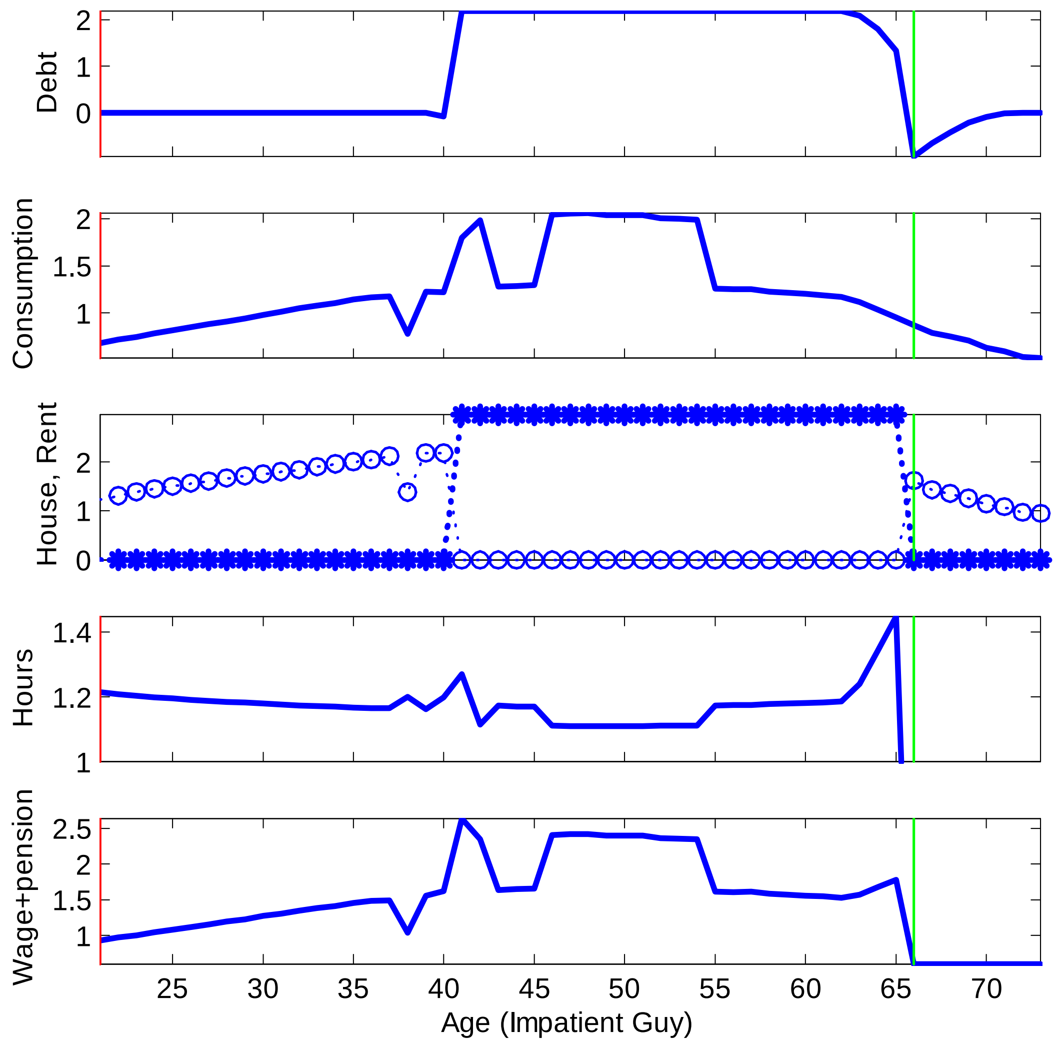

Figure 3: Life-cycle profiles (impatient agent).

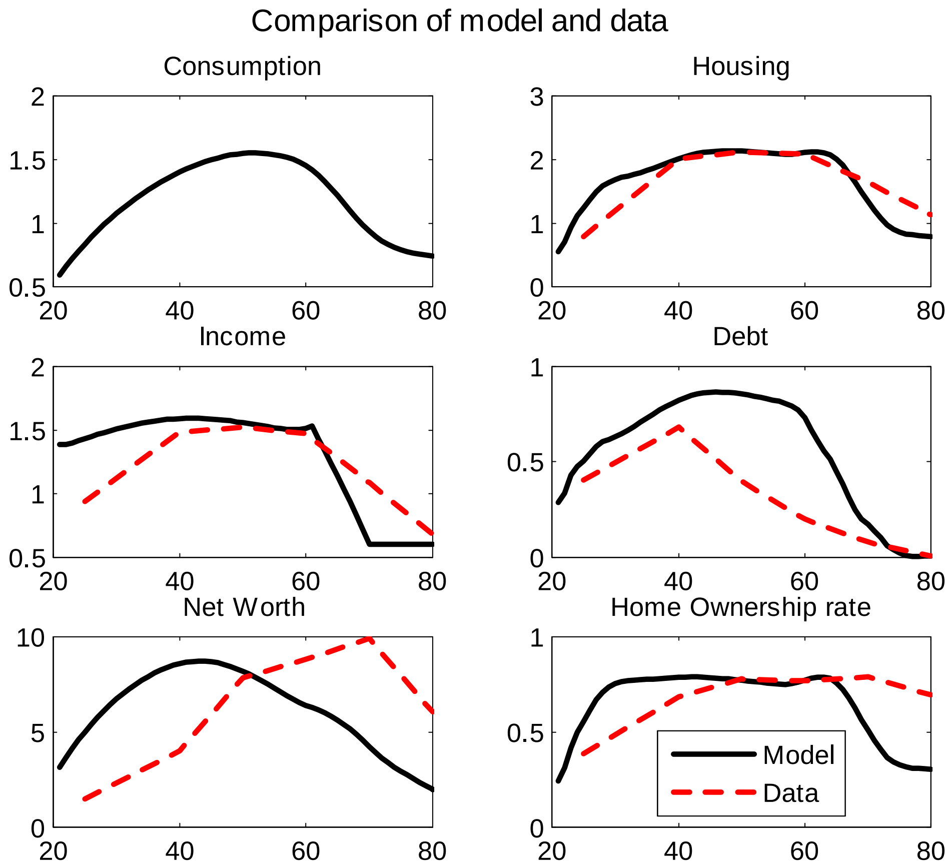

Figure 4: Model (calibration 1) vs data (1983 SCF)

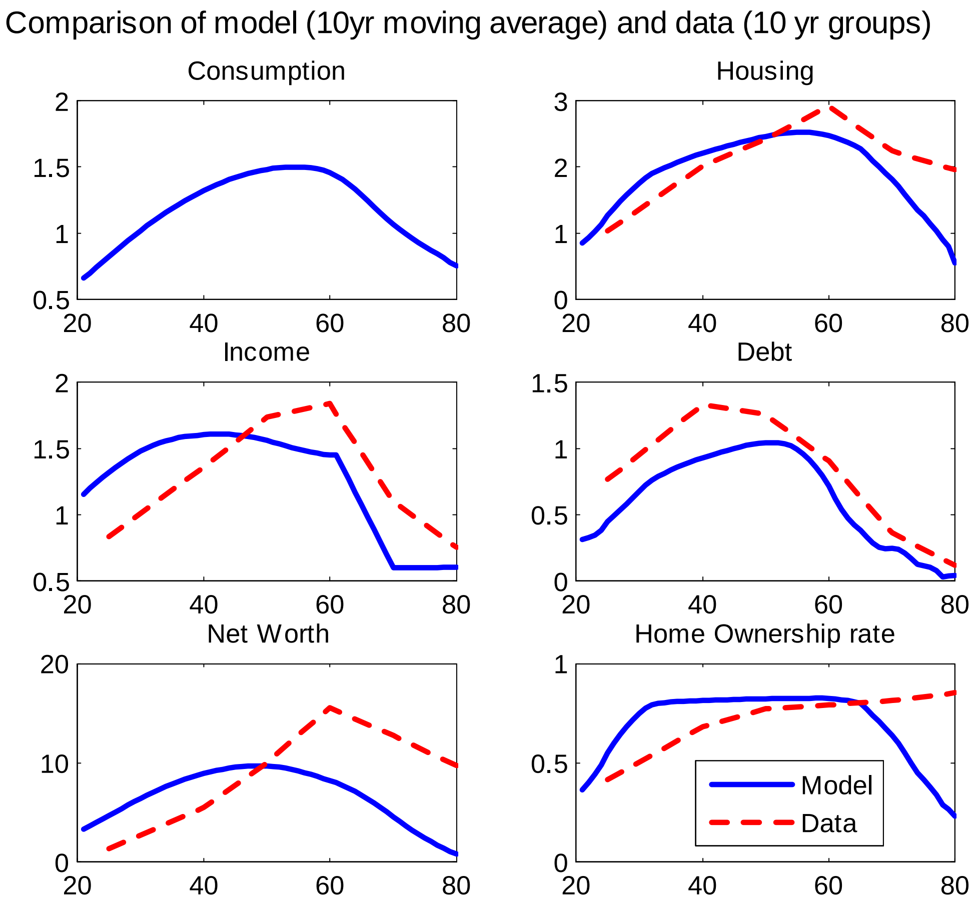

Figure 5: Model (calibration 4) vs data (2004 SCF)

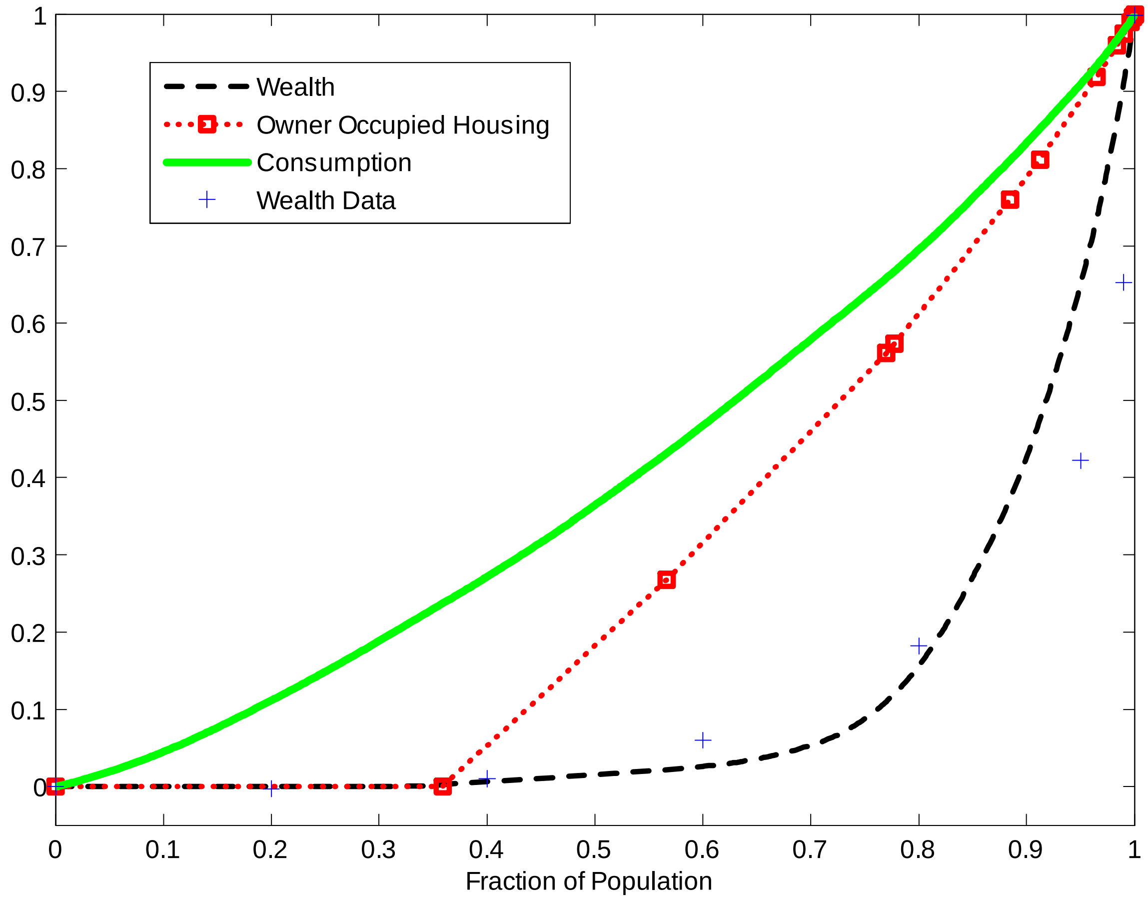

Figure 6: Lorenz curves for wealth, owner-occupied housing and consumption in the model (baseline calibration). The + sign refers to the data (source Budria Rodriguez et al, 2002).

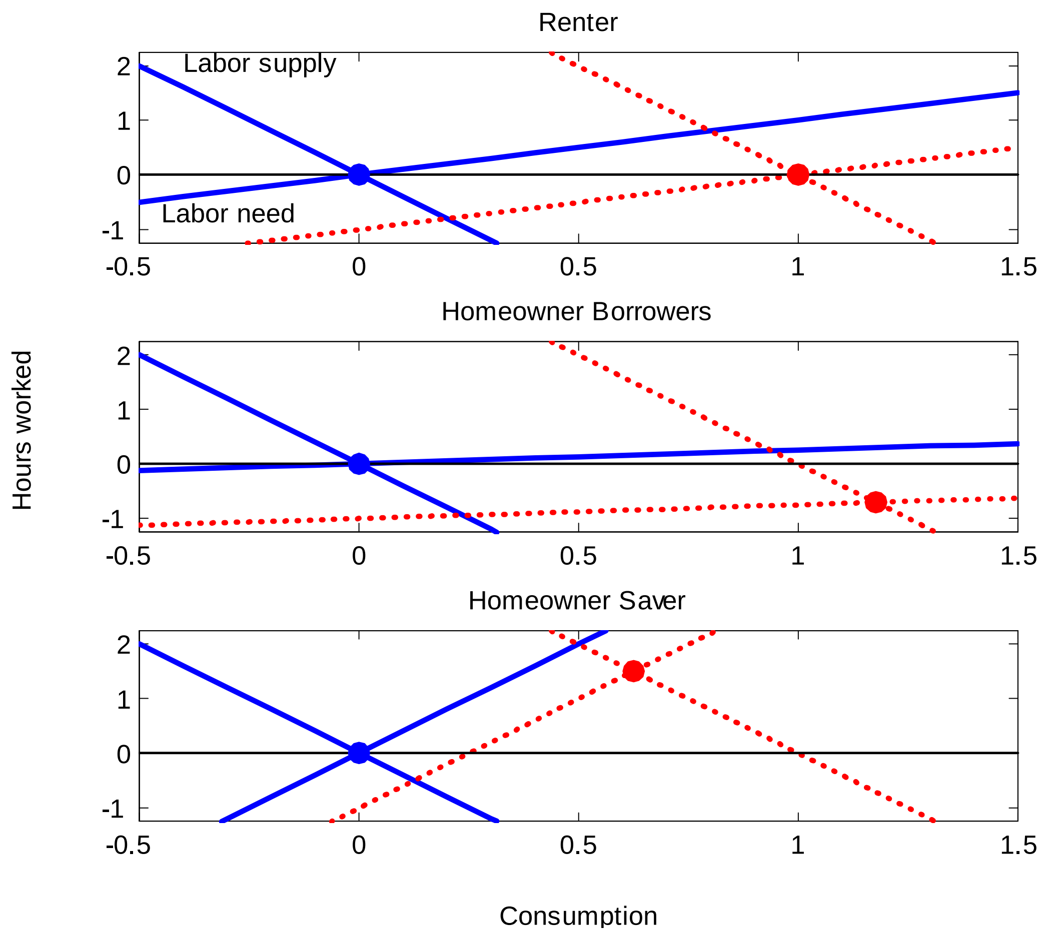

Figure 7: Equilibrium hours worked for renters, borrowers and savers in response to wage changes. The horizontal and vertical axis plot respectively percentage deviations of consumption and hours from their steady state values. The downward sloping line plots is the labor supply curve as a function of consumption (the negative slope reflects the negative wealth effect on labor supply from higher consumption). The upward sloping line is the labor need curve from the household budget constraint (the positive slope reflects the need to work more to finance higher consumption needs). Increases in the wage move both lines to the right.

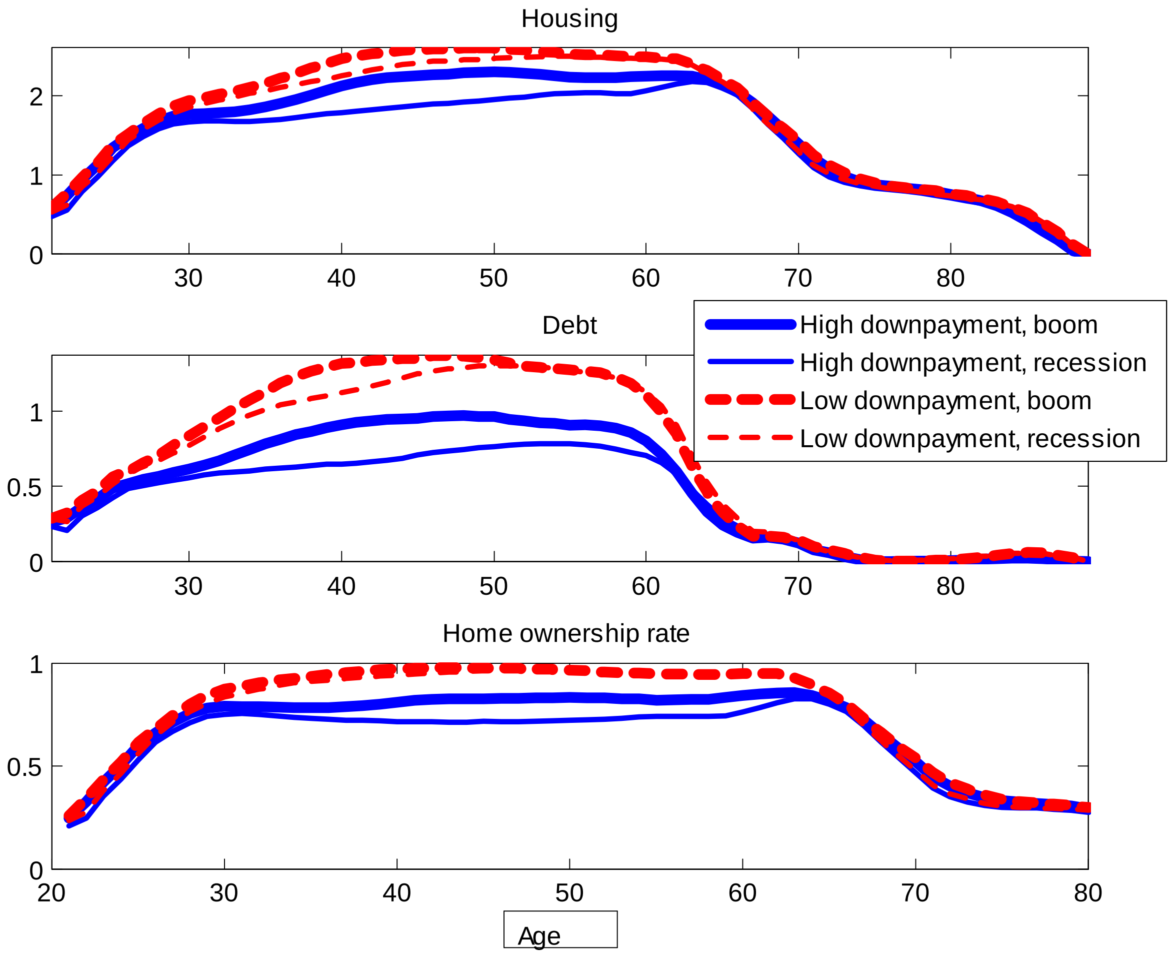

Figure 8: Business cycle implications of high (blue, solid) and low down-payment requirements (red, dashed lines) for housing and home ownership rates. With low downpayment requirements, housing investment is less volatile, especially among middle-aged people.