Working Papers

FEDERAL RESERVE BANK OF BOSTON™

No. 09-12

Housing and Debt Over the Life Cycle and Over the Business Cycle

Matteo Iacoviello and Marina Pavan

Abstract:

This paper describes an equilibrium life-cycle model of housing where nonconvex adjustment costs lead households to adjust their housing choice infrequently and by large amounts when they do so. 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.

The authors use a calibrated version of their model to ask the following question: what are the consequences for aggregate volatility of an increase in household income and a decrease in downpayment requirements? They distinguish between an early period, the 1950s though the 1970s, when household income risk was relatively small and loan-to-value ratios were low, and a late period, the 1980s through today, with high household income risk and high loan-to-value ratios. In the early period, precautionary saving is small, wealth-poor people are close to their maximum borrowing limit, and housing investment, homeownership, and household 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 found in the data. Quantitatively, this model can explain: (1) 45 percent of the reduction in the volatility of household investment; (2) the decline in the correlation between household debt and economic activity; and (3) about 10 percent of the reduction in the volatility of GDP.

JEL Classifications: E22, E32, E44, E51, D92, R21

Matteo Iacoviello is an associate professor of economics at Boston College and a visiting scholar at the Federal Reserve Bank of Boston; his e-mail address is iacoviel@bc.edu. Marina Pavan is a senior research fellow at University College Dublin's Geary Institute; her e-mail address is marina.pavan@ucd.ie.

This paper, which may be revised, is available on the web site of the Federal Reserve Bank of Boston at http://www.bos.frb.org/economic/wp/index.htm.

This paper was written while Matteo Iacoviello was a visiting scholar at the Federal Reserve Bank of Boston: we are grateful for the hospitality. We also thank Massimo Giovannini and Joachim Goeschel for their invaluable research assistance. We thank Chris Carroll, Kalin Nikolov, Dirk Krueger, and Makoto Nakajima, as well as various seminar and conference participants for helpful comments on various drafts of this paper.

The views and opinions expressed in this paper are those of the authors and do not necessarily represent those of the Federal Reserve Bank of Boston or the Federal Reserve System.

This version: November 2, 2009

1. Introduction

Housing investment is a volatile component of GDP.[1] Historically, this observation has led researchers to emphasize movements in the housing market as central to understanding aggregate fluctuations. However, modern business cycle theory has often been silent on this topic. When housing is included in dynamic stochastic equilibrium models, its role is inconsistent with its definition: there is no role for income and wealth heterogeneity, no borrowing constraints, no distinction between owning and renting, no transaction costs (or unrealistic ones) for adjusting home size, and no life-cycle considerations.[2]

Our goal in this paper is to address this imbalance. Specifically, 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[3] 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 productivity time-varying. We introduce housing by modeling some key features that make housing different from other goods: its role as collateral for loans, its lumpiness, and the choice of renting versus owning. Finally, we relax the assumption that households have identical tastes by splitting the population in one "patient" and one "impatient" group: this simple modification makes the wealth distribution highly skewed, in a manner similar to the data.[4]

Results. 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 well replicates the life-cycle profiles of housing and nonhousing wealth. The young, the old and the poor become renters and hold few assets. The middle-aged and the wealth-rich become homeowners. For a typical household, the asset portfolio is simple: it consists of a house and a large mortgage. Despite its stylized nature, the model also reproduces the frequency and size of individual housing adjustment: homeowners change house size infrequently and in large amounts when they do so; renters change house size often, but in smaller amounts.

In terms of its business cycle properties, our model replicates two empirical characteristics of housing investment: its procyclicality and its high volatility. In addition, the model is able to match 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

[1] Throughout the paper, we use the terms "housing investment", "household investment," and "residential investment" interchangeably. Their data counterpart in the National Income and Products Accounts (unless otherwise noted) is Real Private Residential Fixed Investment.

[2] See Davis and Heathcote (2005) and Fisher (2007) for examples of equilibrium models with housing.

[3] In the paper, we use the terms idiosyncratic risk, household income risk, and individual income risk interchangeably.

[4] 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.

quantitative general equilibrium.

Model Experiments. We illustrate the workings of our model with two experiments characterizing the business cycle implications of (1) increasing micro volatility (household income risk) and (2) lowering downpayment constraints.[5] These structural changes, which occurred roughly around the 1980s, might have affected the sensitivity of macroeconomic aggregates to given economic shocks, and are potential candidates for explaining the role of debt and the housing market in the Great Moderation, especially given two observations on the post-1980s U.S. economy (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 0.29.[6] We single out these two changes because we regard risk and availability of finance as two key determinants of housing demand and housing tenure: higher risk should make individuals more reluctant to buy large items that are costly to sell in bad times; greater available financing should encourage housing demand, since it would reduce the amount of savings that are necessary to buy a given size house.

In line with the data, we find that the combination of larger idiosyncratic risk and lower downpayment requirements can (1) reduce the relative volatility of housing investment; and (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 (relative to renters) are more likely to work more in bad times in order to finance housing and mortgage payments, thus offsetting the decrease in output due to negative productivity 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 home buyers). Both these forces reduce housing investment volatility and aggregate volatility.

An increase in idiosyncratic risk leads to higher precautionary saving, and to a decrease in homeownership rates among impatient agents. In addition, larger risk makes wealth-poor individuals more cautious: that is, wealth-poor individuals adjust their consumption, hours and housing demand by smaller amounts in response to aggregate shocks. This mechanism is particularly pronounced for housing purchases, since a house is a large item that is costly to purchase and sell. Combined with low downpayment requirements, this effect reduces the procyclicality of household debt and reduces the sensitivity of housing demand to given changes in aggregate conditions.

We find that the changes in income volatility and in downpayment requirements help form a

[5] See, for instance, Campbell and Hercowitz (2005) and Gerardi, Rosen and Willen (forthcoming) for the role of financial reforms, and Dynan, Elmendorf and Sichel (2007) for a discussion on the evolution of household income volatility.

[6] See Stock and Watson (2002) for an overview of the research on the Great Moderation.

qualitative and quantitative explanation of the changes observed in the data. Together, they can explain more than 10 percent in the reduction in the variance of GDP, 45 percent of the reduction in the variance of housing investment, and the entire decline in the correlation between household 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. However, most of this literature abstracts from housing altogether or implicitly considers housing as part of the total capital stock.[7] Some exceptions are discussed below.

Silos (2007) analyzes the relationship between macroeconomic shocks and household portfolio choice by adopting a life-cycle framework, but does not model the extensive margin of owning versus renting and assumes convex costs for making housing adjustment.[8] 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 recent 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. Their approach, however, sidesteps general equilibrium considerations since they keep the interest rate constant. Kiyotaki, Michaelides and Nikolov (2007) use a stylized life-cycle model of housing tenure to study the interaction between borrowing constraints, house and land prices, and economic activity. Favilukis, Ludvigson and Van Nieuwerburgh (2009) use a two-sector RBC model with housing that also considers the interaction between borrowing constraints and aggregate economic activity, but address a different set of questions than we do. Finally, closely related to our approach is the paper by Campbell and Hercowitz (2005): they study the impact of financial innovation on macroeconomic volatility in an infinite horizon 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 important features with ours.[9]

[7] Papers on housing in incomplete market models with heterogeneous agents that abstract from aggregate shocks include Gervais (2002), Fernandez-Villaverde and Krueger (2004), Ortalo-Magné and Rady (2006) and Gruber and Martin (2003), among others.

[8] Convex adjustment costs for housing induce adjustment dynamics that are dramatically 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. See also Carroll and Dunn (1997) for an early partial equilibrium model with (S, s) behavior for housing.

[9] Nakajima (2005) uses an incomplete markets model with a fixed housing supply to study the relationship between

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 Ríos-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 households, 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, he is replaced by a working age descendant who inherits the dead person's estate and earnings process.

Denote with χ_{a+1} the probability of surviving from age a to $a+1$, and let Π_a be the stationary distribution of individuals 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: $\Pi_1 = \sum_{a=1}^{T}(1-\chi_{a+1})\Pi_a$. Let the measure of all individuals at any given period be normalized to one: $\sum_{a=1}^{T}\Pi_a = 1$.

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

1. Their age;
2. Their labor productivity;
3. Their patience. A recent literature suggests that preference heterogeneity may be an important source of wealth inequality. This idea 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 onset 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.[10]

Households receive utility from the nondurable consumption, housing services and leisure. They can choose between renting housing services and owning housing capital. There are no state contingent markets for hedging against idiosyncratic risk, and the only self-insurance possible is through claims on the economywide capital stock. 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

earnings inequality and housing prices. Van Nieuwerburgh and Weill (2006) study the joint dynamics of income and housing prices and quantities to study the quantitative effects of income inequality and housing supply regulation on the dispersion of house prices at the regional level.

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

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 nonhousing investment over the life cycle and the business cycle.

2.1. Household Preferences and Endowments

Let $\bar{l}$ denote each agent's total time endowment. Households derive utility from leisure $(\bar{l} - l)$, nondurable 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 simple formulation:

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

Above j and τ are positive, and $\theta = 1$ if $s = h > 0$ (meaning the individual owns his house), while $\theta < 1$ if $h = 0$ (meaning the individual rents). The assumption for θ implies that a household experiences a net utility gain when transitioning from renting to owning a home and is standard in models of homeownership in the public economics and urban economics literature; see Rosen (1985) and Poterba (1992). We also assume that, when individuals are homeowners, there is a minimum size house $\underline{h}$ that can be purchased, and that rental units may come in smaller sizes than houses, thus allowing renters to consume a 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 the 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 the 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. 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. The total net income at age a in period t is denoted by y_{at}. Then:

$$y_{at} = w_t \eta_a z_t l_t - \Gamma \text{ if } a \leq \widetilde{T}, \tag{2}$$

$$y_{at} = P \text{ if } a > \widetilde{T}. \tag{3}$$

Households start their life with endowments given by b_0 and h_0, the accidental bequests left by a deceased agent.[11] Households can buy and sell only one bond, b, which pays a gross interest rate of

[11] We assume that individuals do not have any motive for leaving intended bequests. Our main results were not sensitive to this modeling assumption.

R_t in period t. Let positive amounts of this bond denote a net debt position.[12] 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 proxy for their expected lifetime earnings:

$$b_t \leq \min\{m_h h_t, m_y \Re_t (y_{at}; R_t, w_t)\}. \tag{4}$$

Above, $\Re_t (y_{at}; R_t, w_t) = y_{at} + \sum_{s=a+1}^{T} \frac{E_t(y_s|y_{at};w_t)}{(R_t)^{s-a}}$ 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).[13] The specification of borrowing constraint (1) rules out unsecured debt; (2) restricts debt to homeowners only; and (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. Finally, we introduce a life-cycle preference shifter λ_a in the utility function. Changes in λ_a mimic changes in household size that deterministically affect the marginal utility of consumption, as in Cagetti (2003). Summing up, households maximize their expected lifetime utility:

$$E_1 \left(\sum_{a=1}^{T} \beta_i^{a-1} \lambda_a (\prod_{\tau=1}^{a-1} \chi_{\tau+1}) u \left(c_a, s_a, \bar{l} - l_a\right) \right) \tag{5}$$

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.

2.2. The Financial Sector and the Housing Rental Market

A perfectly competitive financial sector collects deposits from households who save, lends to firms and households who borrow, 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.[14] This technological assumption guarantees that the prices of housing and of capital in units of consumption are constant. Let p_t be the price of each unit of rental services at time t. Then a no-arbitrage condition holds such

[12] We therefore refer to b as financial liabilities (or net debt), and correspondingly to $-b$ as financial assets (or net assets). Because bonds are claims on aggregate capital, their return varies with the aggregate state.

[13] The measure is approximated since the interest and wage rates are fixed at the current values, and because the approximation does not account for the endogeneity of labor supply (when computing y_{at}, we assume $l_t = \bar{l}$ at any $t = 1, .., \widetilde{T}$). This greatly simplifies the calculation of the expected lifetime earnings, allowing us to derive a constraint which prevents the elderly from borrowing too much.

[14] 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 sector finances itself by issuing loans in t that pay a gross interest rate R_{t+1} in period $t + 1$.

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,

$$p_t = E_t \left(\frac{R_{t+1} - (1 - \delta_H)}{R_{t+1}} \right) \tag{6}$$

at any t, where δ_H is the depreciation rate of the housing stock.[15]

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 . Output is produced according to the Cobb-Douglas technology,

$$Y_t = A K_{t-1}^{\alpha} L_t^{1-\alpha}, \tag{7}$$

where L and K denote aggregate labor and aggregate capital respectively, $\alpha \in (0, 1)$ is the capital share of aggregate income, while $A \in \widetilde{A} \equiv \left\{A^1, .., A^{n_a}\right\}$ represents a stochastic shock to total factor productivity. This 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 $H_t = H_t^o + H_t^r$ (owned and rented), and the total transaction costs incurred by homeowners for adjustments to the housing stock, which we denote by Ω_t:

$$C_t + H_t - (1 - \delta_H) H_{t-1} + \Omega_t + K_t - (1 - \delta_K) K_{t-1} = Y_t \tag{8}$$

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

Since the economy is closed, 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 $\Phi_t(z_t, b_{t-1}, h_{t-1}; \beta, a)$ the distribution of households over earnings shocks, asset holdings, housing wealth, discount factors and ages in period t. Without aggregate uncertainty, the economy

[15] The expectation term in the no-arbitrage condition reflects the assumption that housing services are paid for and yield utility in the current period, whereas capital services are paid for and yield output with a one-period delay (this timing assumption is standard in decentralized version of the neoclassical growth model). One can interpret the marginal cost of one house to be 1 for the financial sector, since loanable funds can be converted into housing costlessly; and the marginal benefit to be the sum of the current rental income, p_t, 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 equation (6).

would be in a stationary equilibrium, with an invariant distribution Φ and constant prices. Yet given aggregate volatility, 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 future wages and interest rates. Both variables depend on future productivity and aggregate capital-labor ratio, which in turn are determined by the overall distribution of individual states. 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 one moment of the distribution Φ is sufficient to forecast future prices.

We write the household optimization problem in recursive formulation. The state variables in period t 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 the period. We assume that agents only use the mean level of nonhousing wealth in order to predict the next period's wage and interest rates. That is, agents observe beginning of period capital K_{t-1} and approximate the evolution of aggregate capital and labor[16] with a linear function that depends on the aggregate shock A_t. Denote $x_t \equiv (z_t, b_{t-1}, h_{t-1}, A_t, K_{t-1})$ the vector collecting individual and aggregate state variables.[17] In recursive form, the dynamic problem of an age a household with discount factor β_i can be stated as follows:

$$V_a(x_t; \beta_i) = \max_{I^h \in \{0,1\}} \{I^h V_a^h(x_t; \beta_i) + \left(1 - I^h\right) V_a^r(x_t; \beta_i)\} \tag{9}$$

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 to buy/own. The value of being a homeowner is the solution to the following problem:

$$V_a^h(x_t; \beta_i) = \max_{c_t, b_t, h_t, l_t} \{\lambda_a u\left(c_t, h_t, \bar{l} - l_t\right) + \beta_i \chi_{a+1} \textstyle\sum_{z_{t+1}, A_{t+1}} \pi_{A_t, A_{t+1}} \pi_{z_t, z_{t+1}} V_{a+1}(x_{t+1}; \beta_i)\} \tag{10}$$

$$\begin{aligned} \text{s.t.} \qquad & c_t + h_t + \Psi(h_t, h_{t-1}) = y_{at} + b_t - R_t b_{t-1} + (1 - \delta_H) h_{t-1}, \\ & b_t \leq \min\{m_h h_t, m_y \Re_t\}, \; c_t \geqslant 0, \; l_t \in \left(0, \bar{l}\right), \\ & K_t = \digamma^K(K_{t-1}, A_t), \; L_t = \digamma^L(K_{t-1}, A_t). \end{aligned}$$

Here $\digamma^K$ and $\digamma^L$ are linear functions in K_{t-1}, whose parameters depend on the aggregate shock A_t, which denote the law of motion of the aggregate state, which agents take as given. The term Ψ in

[16] In our definition of equilibrium, how aggregate labor is supplied is needed as an input in the agent's decision since wages and interest rates depend on aggregate capital and labor.

[17] We present the "approximate" recursive formulation in which the aggregate state variables are represented by the economy's capital 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.

the budget constraint shows 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 practices 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.[18]

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

$$V_a^r(x_t;\beta_i) = \max_{c_t,b_t,s_t,l_t} \{\lambda_a u(c_t, s_t, \bar{l} - l_t) + \beta_i \chi_{a+1} \textstyle\sum_{z_{t+1},A_{t+1}} \pi_{A_t,A_{t+1}} \pi_{z_t,z_{t+1}} V_{a+1}(x_{t+1};\beta_i)\} \tag{11}$$

$$\begin{aligned} \text{s.t.} \qquad & c_t + p_t s_t + \Psi(0, h_{t-1}) = y_{at} + b_t - R_t b_{t-1} + (1-\delta_H) h_{t-1}, \\ & b_t \leq 0,\ c_t \geqslant 0,\ l_t \in (0,\bar{l}), h_t = 0, \\ & K_t = \digamma^K(K_{t-1}, A_t),\ L_t = \digamma^L(K_{t-1}, A_t). \end{aligned}$$

At the agent's last age, $V_{T+1}(x_{T+1};\beta) = 0$ for any $(x_{T+1};\beta)$.

We are now ready to 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* $\{p_t\}_{t=1}^{\infty}$, *aggregate variables* K_t, L_t, H_t^o *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, *and the laws of motion* $\digamma^K$ *and* $\digamma^L$, *the value functions are the solution to the individual's problem, with the corresponding policy functions.*

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

$$R_t - 1 + \delta_K = \alpha A_t (K_{t-1}/L_t)^{\alpha-1}, \tag{12}$$

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

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

$$p_t = E_t \left(\frac{R_{t+1} - (1-\delta_H)}{R_{t+1}} \right). \tag{14}$$

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

$$-\int b_a(x_t;\beta)\, d\Phi_t(z_t, b_{t-1}, h_{t-1}; \beta, a) + p_t H_t^r = K_t + H_t^r, \tag{15}$$

[18] Implicit in this formulation is the simplifying assumption that the household can adjust the level of housing consumption only by selling the old house and buying a 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, given that the household can expand, remodel, or fail to maintain the house. For simplicity, here we rule this possibility out.

where[19] H_t^r *is the rented housing stock, defined as:*

$$H_t^r = \int (1 - I_a^h(x_t;\beta)) s_a(x_t;\beta)\, d\Phi_t(z_t, b_{t-1}, h_{t-1}; \beta, a). \tag{16}$$

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

$$L_t = \int l_a(x_t;\beta)\, \eta_a z_t d\Phi_t(z_t, b_{t-1}, h_{t-1}; \beta, a) \tag{17}$$

and as a consequence the goods market satisfies the resource feasibility constraint (8), where H_t *and* Ω_t *are defined as follows:*

$$H_t = H_t^o + H_t^r = \int I_a^h(x_t;\beta)\, h_a(x_t;\beta)\, d\Phi_t(z_t, b_{t-1}, h_{t-1}; \beta, a) + H_t^r, \tag{18}$$

$$\Omega_t = \int \Psi(h_a(x_t;\beta), h_{t-1})\, d\Phi_t(z_t, b_{t-1}, h_{t-1}; \beta, a). \tag{19}$$

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

$$\sum\nolimits_{a=1}^{\widetilde{T}} \Pi_a \Gamma = \sum\nolimits_{a=\widetilde{T}+1}^{T} \Pi_a P. \tag{20}$$

6. ***The laws of motion for the aggregate capital and aggregate labor*** *are given by*

$$K_t = \digamma^K(K_{t-1}, A_t), \; L_t = \digamma^L(K_{t-1}, A_t). \tag{21}$$

Appendix A provides the details on our computational strategy.

3. Parameterization

In our baseline calibration, we aim at reproducing basic facts for the U.S. economy from 1952 to 1982. We characterize these three decades as a period of high aggregate volatility, low idiosyncratic volatility, and high downpayment requirements. Later, we will change idiosyncratic volatility and downpayment requirements in order to isolate the role they play in affecting the properties of economic aggregates. Our calibration is summarized in table 2.

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

[19] The rented housing stock is rented out and paid for within the period. For this reason, the market clearing condition requires that total resources available for production of new capital (the sum of K_t and H_t^r) equal supply of net financial assets (that is, $-\int b_a(\cdot)\, d\Phi_t(\cdot)$) plus the flow rental income from the rented housing stock ($p_t H_t^r$).

equal to the survival probabilities for men aged 21-90, taken from the U.S. Decennial Life Tables for 1989-1991.[20] 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, the total population remains constant. Finally, 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 polynomial. 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 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 40 percent of the average labor income in the economy, which is financed through lump-sum taxes.[21] Our idiosyncratic shock to labor productivity is specified as an order-one autoregressive process as follows:

$$\log z_t = \rho_Z \log z_{t-1} + \sigma_Z \left(1 - \rho_Z^2\right)^{1/2} \varepsilon_t, \qquad \varepsilon_t \sim Normal\left(0, 1\right), \tag{22}$$

which we approximate with a three-state Markov process.[22] We set $\rho_Z = 0.9$. In the baseline calibration we set $\sigma_Z = 0.30$. Later, we increase this number to $\sigma_Z = 0.45$ to capture the increased earnings volatility of the 1990s. The calibration of these numbers, which are in the ballpark of the microeconometric estimates, is discussed in more detail in Appendix B.

3.3. Preferences and Housing Adjustment Costs

We assume that there are two classes of households, a "patient" group with a discount factor of 0.995 and an "impatient" group with a discount factor of 0.925. The high discount factor pins the real interest rate down to 3.75 percent. The low discount factor is in the range of estimates in the literature – see, for instance, Hendricks (2007) and references therein. We further assume that one-third of the population is composed of patient households, and the remaining two-thirds of households are impatient. The shares of patient and impatient agents imply that one-third of the population hold most of wealth, and deliver a Gini coefficient for wealth around 0.75, in line with the data. In the robustness section, we discuss the unappealing properties of the model that assumes that all people have identical tastes. We set $\tau = 1.5$ and the total endowment of time $\bar{l} = 2.5$; these

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

[21] Queisser and Whitehouse (2005) report that average pensions for males in the United States are 40 percent of the economy-wide average earnings.

[22] We do so using the procedure described in Tauchen (1986).

two parameters imply that the time spent working in the market (around 1.1) is 40 percent of the available agents' 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 average housing investment to private output ratios around 6 percent, and a ratio of the housing stock to output 1.22. 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 low-range estimate of the actual costs of moving and changing a house: our model does not allow small adjustments to housing consumption (such as improvements and failure to maintain), so in absence of this margin we choose to be conservative on this value.[23]

3.4. Technology and Borrowing Constraints

We set the capital share $\alpha = 0.33$ and the capital depreciation rate $\delta_K = 0.10$. In the economies that we consider, these values yield average capital to output ratios around 2.4 and average business investment to output ratios around 25 percent on an annual basis.[24]

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 seven states to match the following first-order autoregressive representation for the logarithm of total factor productivity:

$$\log A_t = \rho_A \log A_{t-1} + \sigma_A \left(1 - \rho_A^2\right)^{1/2} \varepsilon_t, \qquad \varepsilon_t \sim Normal\,(0, 1)\,. \tag{23}$$

We set $\rho_A = 0.925$ and $\sigma_A = 0.0139$. After rounding, the first number mimics 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 output to that of the detrended data.

Last, in the baseline calibration we set the maximum loan-to-value ratio m_h at 0.75. We increase this number to 0.85 in the calibration for 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: with these numbers, the income constraint is binding only late in life, and essentially prevents older homeowners from borrowing. 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.

[23] Larger, probably equally plausible adjustment costs (say, $\psi = 8\%$) imply very infrequent changes in housing size for homeowners (the median owner changes house less than once during his lifetime).

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

3.5. Homeownership and Renting

We assume that the minimum-size house size available for purchase ($\underline{h}$) costs two times the average annual pre-tax household income: this constraint – together with the downpayment constraint – prevents asset-poor people from becoming homeowners too quickly.[25] We do not impose any constraint on the size of the rental properties. 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 (θ). 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 $\theta = 0.832$. This number implies that an agent is indifferent between renting a house of $1,000$ square feet and owning a house of 832 square feet.[26]

4. Results

We present 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

General Features of Household Behavior. At each stage in its life-cycle, the household chooses its consumption, saving, labor supply, and housing investment by taking into account current and expected income, and its liquid assets and housing position beginning of period. 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 illustrating the household's dynamic labor supply behavior to the next section, when we discuss the model's dynamics in response to aggregate shocks.

It is simple to characterize the behavior of agents by conditioning on whether they enter the period as renters or homeowners. For renters, the housing investment decision is as follows: at any given age and for any initial state, there is a threshold amount of liquid assets ($-b$ in the notation of the model) such that, if assets exceeds the threshold, renters become homeowners. The larger its initial liquid assets are, the less likely a household is to borrow against its housing purchase.

For existing homeowners, there are four possibilities: homeowners can stay put, increase their house size, downsize or switch to renting. Which option they choose depends on the combination of housing and liquid assets they possess when they enter the period, as well as on age and income.

[25] In the United States, the median house price has averaged around three times the median household income.

[26] 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 downpayment (these outcomes are an equilibrium feature of the model when we set $\theta = 1$).

Figure 2 plots the optimal housing choice for a homeowner as a function of its initial house size and liquid wealth.[27] The solid, downward sloping line plots the borrowing constraint that restricts debt from exceeding a fraction m_h of its 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 keeps its housing stock constant. If its 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-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 its 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.[28]

Life-Cycle Profiles. Figure 3 plots an example of the typical life-cycle choices of an agent in our model. We choose an agent with a low discount factor since the behavior of an agent with low assets and who is close to the borrowing constraint is illustrative of the main workings of the model. This agent starts adult life – at age 21 – with zero assets and a low idiosyncratic income realization and becomes a renter. Over time, as his income rises because of life-cycle reasons, he consumes more and works less. At the age of 36 the agent is hit by a positive income shock, saves just enough to afford the downpayment (his debt becomes negative for one period) and buys a house. The decision to buy a house goes hand in hand with the decision to work more. A positive and temporary income shock raises the incentive to work; moreover, the incentive to work more is stronger for those who pay the fixed cost and buy the house, since these individuals need to set resources aside to afford to downpayment. While he is a homeowner, hours worked move in opposite direction to the wage shocks, rising in bad times, falling in good times (this mechanism is explained in detail in section 4.2 below). When the agent turns 50 years old, he becomes a renter again (after a series of bad income shocks). He enjoys one more spell of homeownership around retirement (which occurs at age 66). Sometime after retirement, the agent switches to a small rental property, and he dies when he turns 78 years old.

An interesting dimension where it is useful to compare the model with the data is the frequency of housing adjustment for homeowners.[29] The work by Hansen (1998), based on the 1993 Survey of

[27] The figure is plotted for a patient agent who is entering retirement (65 years old), when aggregate productivity and the capital-labor ratio are equal to their average value.

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

[29] 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.

Income and Program Participation, reports that the median homeowner stays in the same house for about eight years. Anily, Hornik, and Israeli (1999) estimate that the average homeowner lives in the same residence for 13 years. The corresponding number for our model is about 15 years, hence in line with the data.[30]

Figure 4 compares the age profiles of housing holdings, debt and homeownership with their data counterparts. Like the data, the model is able to capture the hump-shaped profiles of these variables. There are, however, two discrepancies: as for mortgage debt, the model slightly underpredicts debt early in life, and overpredicts debt late in life. The model also tends to underpredict home ownership late in life.[31]

The Wealth Distribution. Our model reproduces the U.S. wealth distribution almost exactly. The Lorenz curves for the U.S. economy and for our model economy are reported in figure 5. The Gini coefficient for wealth in the model is 0.76, and is about the same as in the data (equal to 0.78, according to Budria et al., 2002). The main discrepancy between the model and the data is that we underestimate the fraction of wealth held by the top 5 percent of the population. On the positive side, the model does well at matching the fraction of wealth held by the poorest 40 percent of the U.S. population, which has no assets and no debt, like the renters in our model. Instead, a model without preference heterogeneity would do much worse: in section 6.1 below, we show that the Gini coefficient for wealth in the model with a single discount factor is 0.52, much lower than in the data.

4.2. Business Cycle Results

In this section, we illustrate 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 have studied these features in life-cycle models with aggregate shocks, we do not devote much space here to exploring these issues.[32] 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 nonhousing assets); in turn, heterogeneity in wealth implies different individual responses to the same shock.

[30] If anything, our model should underpredict the frequency of housing adjustment, since it abstracts from "moving shocks" that force homeowners to sell and buy a house even though their housing consumption remains constant. In the data, 15 percent of the moves are associated with a move to a different state, and 35 percent of the moves are associated with a move to a different county. Most of these moves probably "moving shocks" rather than movements along the housing ladder.

[31] Recall that the model is calibrated to target average homeownership rates, so as a consequence the model overpredicts ownership early in life.

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

Workings of the Model. We focus on the response of aggregate hours to a technology shock, since movements in hours are the key element of the propagation mechanism in models that rely on technology shocks are sources of aggregate fluctuations. In particular, we study how the wealth distribution and its composition affect agents' responses to shocks.[33] To fix ideas, consider a stripped-down version of the budget constraint of a working individual that keeps wealth constant between two periods, so that $b_t = b_{t-1}$ and $h_t = h_{t-1}$.[34] Abstracting from taxes and pensions, this implies the following budget constraint:

$$c_t = \underbrace{w_t \eta_a z_t l_t}_{\text{wage income}} + \underbrace{\xi_t}_{\text{interest income net of housing maintenance}} \tag{24}$$

where $\xi_t = -(R_t - 1)\, b_{t-1} - \delta_H h_{t-1}$ measures the resources besides wage income that can be used to finance consumption:[35] the term $(1 - R)\, b$ denotes net interest income; the term $\delta_H h$ denotes the maintenance cost that is required to keep housing 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 (negative b), ξ is positive and large, and wage income is a small fraction of nondurable consumption c. For a renter, h equals 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,[36] 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:

$$\widehat{c}_t = \frac{wl}{c}\left(\widehat{w}_t + \widehat{l}_t\right) \text{ or } \widehat{l}_t = -\widehat{w}_t + \frac{c}{wl}\widehat{c}_t. \tag{25}$$

This version of the constraint can be interpreted, for given 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 change in hours $\widehat{l}$ needed to finance the consumption change, with an elasticity of hours to consumption

[33] Domeij and Floden (2006) use a model with a zero-borrowing constraint to make the point that labor supply estimates are biased downward if liquidity constraints are ignored, and show that labor supply curves can bend backwards for low levels of wealth. Ziliak and Kniesner (1999) find that male labor supply elasticities increase with wealth.

[34] Obviously, the household's 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, we can study the joint determination of consumption and hours by focusing on the budget constraint and the Euler equation for labor supply only. This is a reasonable assumption for small shocks (such as aggregate shocks).

[35] For a renter, housing and nonhousing expenditures are proportional to each other, so the constraint reads as

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

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

[36] This assumption is approximately valid to the extent that movements in the interest rate are not too large, since b and h are state variables from time t's perspective

given by consumption–wage income ratio $\frac{c}{wl} \equiv \phi$. 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, $\phi = 1$; finally, 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 words, a wealthy person needs to increase hours worked by more than 1 percent to finance a 1 percent rise in desired consumption, since labor income makes for less than 100 percent of its consumption; an indebted homeowner needs to increase hours worked by less than 1 percent to finance a 1 percent rise in consumption, because of the leverage effect; a renter needs to increase hours 1 for 1 with consumption. These relationships are plotted in figure 6 for the three types. In the consumption-hours space, the labor need curve is upward sloping, with a slope given by the consumption/wage ratio ϕ; and it is shifted 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

$$\widehat{l} = \zeta \left(\widehat{w} - \widehat{c} \right) \tag{26}$$

where ζ is the steady-state Frisch labor supply elasticity. This curve slopes downward (in the consumption-hours space) because of the wealth effect on labor, and is shifted to the right by a rise in the wage.

Figure 6 also shows how the "labor supply" and "labor need" curves move for a given change in the wage rate, 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 percent: from the standpoint of the budget constraint, the individual can keep constant consumption by reducing hours in proportion to the increase in the wage; it also shifts the labor supply curve up by ζ percent. Which effect dominates depends on whether the consumption–wage ratio ϕ is equal or smaller than one: 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.[37]

For the economy as a whole, the total response of hours to a wage change will be an average of the labor supply responses of all households. If individual labor supply schedules were linear in net wealth, the aggregate labor supply response would be linear in wealth too, and there would be no effect of the wealth distribution on the aggregate labor supply elasticity. There are, however, two main forces that undo this linearity in our setup. First, retirees do not work, so any transfer of wealth to and from this group could affect how the workers respond to wage shocks. For instance, if the return on wealth rises in good times, this reduces the net worth of wealth-poor people, and reduces the incentive to work for those in the labor force, all else equal. Second, the interaction between borrowing constraints and housing purchases creates an interesting nonlinearity. Above, we have assumed that households do not change wealth in response to a shock in the wage. However, if

[37] Combining (25) and (26) and solving for hours as a function of the wage yields $\widehat{l} = \zeta \left(\frac{\phi - 1}{\zeta + \phi} \right) \widehat{w}$.

households switch from renting to owning (or if they increase their house size) in good times, they typically need to save for the downpayment. To do so, they set aside part of their current wage income and cut back in consumption. The "forced savings" at the time of the purchase of the house shifts the labor need curve to the left: intuitively, if the individual wants to keep consumption constant when he buys the house, he needs to work even more hours. This effect creates an important comovement between hours worked, on one hand, and housing purchases, on the other.[38] 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. 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 HP-filtered, post-war U.S. data, the relative variability of housing investment is large, with a standard deviation that is between three and four 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 less procyclical thereafter, with a correlation with GDP that drops from 0.78 to 0.29.

Against this data background, table 3 reports some of the key statistics generated by the benchmark model, and compares them to the data. Overall, our baseline model does 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 percent of the total variance 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 the data, but, as in the data, it is smaller than the volatility of housing investment: this result occurs 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 its cyclical behavior.[39] The key to this result is the fact the bulk of the population (the impatient two thirds) upgrades its housing in good times by taking out a mortgage at the same time; in particular, these are very young and the middle-aged. At the same time, the model tends to overpredict the volatility of debt itself: the standard deviation of the model variable is about four times larger than in the data. We suspect that one possible reason for the higher volatility of the model variable has to do with the

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

[39] We define household debt in the model as $D_t = \int_{b>0} b_a(x_t; \beta)\, d\Phi_t$ (that is, the average of the household liabilities).

simplifying assumption of our model that only one financial asset is available, whereas in the data some households (especially the wealthy) own simultaneoulsy 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.

A dimension where it is interesting to compare the model with the data pertains to home sales. In our model, we classify as a home sale every instance in which a household pays a transaction cost to change its housing choice: this involves own-to-own, rent-to-own and own-to-rent transitions. By this metric, in every period between 3 and 4 percent of the model's housing stock changes hands. The model correlation between (log) output and sales is positive, about 0.8, and the standard deviation of sales is 0.4 percent. These numbers are very much in line with the data. In the 1952-1982 period, existing home sales are between 3 and 4 percent of the total housing stock, the correlation between sales and GDP is 0.7, and the standard deviation of sales is 0.43 percent. The positive correlation between sales and economic activity that the model captures reflects the presence of borrowing constraints: when the economy is in recession and household balance sheets are 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 to do so.[40]

5. Dynamic Consequences of Higher Idiosyncratic Volatility and Lower Downpayments

Having shown above that the model roughly captures postwar U.S. business cycles, we now consider the implications of two experiments. In the first, we lower the downpayment requirement from 25 to 15 percent of the purchase price. In the second, we increase the amount of idiosyncratic risk faced by households, changing the unconditional standard deviation of income σ_Z from 0.30 to 0.45. 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 that began in the mid 1980s. The results are shown in table 4.

5.1. A Decrease in Downpayment Requirements

A larger value of m_h (that is, a decrease in the required downpayment $1 - m_h$) has two main effects on the properties of the model. First, it leads to an increase in the homeownership rate. Second, it reduces the volatility of household investment and, to a lesser extent, of the other components of demand.

Lower required downpayments allow more housing ownership among the portion of the population

[40] In the model with a single discount factor and very few wealth-poor people the correlation between GDP and sales becomes much smaller (it equals 0.25), because this mechanism is absent.

with very little net worth. In particular, the model's prediction is that lower downpayments substantially increase homeownership for those between the ages of 30 and 65 years. The homeownership rate rises from 64 to 76 percent.

Turning to business cycles, the rise in m_h reduces the volatility of housing investment, from about 6.69 to 6.41 percent. Why? There are three main forces at work. One works directly though a "demand-side" channel. When downpayment requirements are high (low m_h), 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. By contrast, when downpayment requirements are low (high m_h), less people are financially 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 downpayment requirement leads to lower volatility of housing investment.

The second force has to do with adjustment costs. On average, because of adjustment costs, homeowners modify their housing consumption little over time relative to renters. This reduces the volatility of housing investment more than the other components of expenditure.

The third force that reduces the volatility of housing investment operates through the the interaction of labor supply and housing purchases. As we explained above, indebted homeowners are more likely, compared to renters, to reduce hours worked in response to positive technology shocks, so their presence dampens aggregate shocks. Therefore, the higher homeownership rate induced by looser borrowing constraints contributes to lower aggregate volatility.[41]

5.2. An Increase in Individual Earnings Volatility

An increase in individual earnings volatility slightly reduces homeownership. The homeownership rate falls from 64 to 62 percent. The lower homeownership rate would tend to increase the volatility of total housing investment, leaving unaffected the other properties of the model. However, higher earnings volatility makes homeowners more reluctant to change their housing consumption: the main mechanism at work is the combination of uncertainty and adjustment costs. When changing housing consumption involves paying transaction costs, existing homeowners are less likely to modify their

[41] A similar intuition has been proposed and analyzed in Campbell and Hercowitz (2005), who show that financial innovation alone can explain a large fraction – more than half – of the reduction in aggregate volatility in a calibrated equilibrium model with borrowers and lenders and downpayment constraints. Aside from modeling differences – our model considers the owning/renting margin and addresses issues related to life cycle, lumpiness and idiosyncratic risk that are absent in their setup –, the intuition they offer for their result carries over to our model, but we find that the effect of lower downpayment requirements is quantitatively smaller. We conjecture that the quantitative differences depend on one modeling assumption: in our setup, the presence indebted homeowners mitigates aggregate volatility, but this effect is partly offset in general equilibrium 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 the labor supply of wealthy homeowners is constant, thus killing this offsetting mechanism.

housing position in response to changes in their net worth. This occurs because modifying housing choice 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 – an issue that we return to in the next subsection. Quantitatively, the higher earnings volatility reduces the standard deviation of housing investment from 6.69 to 6.3 percent.

5.3. Combining Lower Downpayments and Higher Volatility

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

The last column of table 4 shows the business cycle consequences of these two structural changes combined: 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 less procyclical, as in the data. The correlation between debt and output falls from 0.82 to 0.35, a change of similar magnitude to the data (from 0.78 to 0.29). Overall, the combination of lower downpayments and high idiosyncratic volatility reduces the standard deviation of GDP from 2.09 percent to 2.02 percent, and the standard deviation of residential investment from 6.69 percent to 5.34 percent.[42] If we interpret these two structural changes as occurring during the 1980s, the corresponding numbers in the data are as follows: the standard deviation of GDP falls from 2.09 percent to 1.27, and the standard deviation of residential investment falls from 8.23 to 3.77 percent. Hence, the two structural changes combined can explain a good size of the reduction in aggregate volatility. Precisely stated, 10 percent of the variance reduction in GDP and 45 percent of the variance reduction in residential investment can be explained by the two factors we highlight here.

Our interpretation of these results is as follows: in response to a combination of high leverage (induced by lower downpayments) and higher income volatility, highly leveraged households become more *cautious* in response to aggregate shocks,[43] thus reducing the extent to which they change their borrowing and their housing demand when aggregate productivity changes. This is especially true for housing – relative to other categories of expenditure – since housing purchases involve durable goods and are subject to substantial adjustment costs. Because individuals are more reluctant to adjust their housing consumption during uncertain times, the sensitivity of hours to aggregate shocks

[42] In experiments not reported here, we find that transaction costs for housing adjustment are crucial to explain the reduction in the volatility of housing investment, since it is their interact with idiosyncratic risk that makes individuals reluctant to adjust their housing consumption in response to aggregate shocks. Without adjustment costs, lower downpayments and higher risk leave the volatility of housing investment and aggregate GDP virtually unaffected.

[43] A cautious individual, according to The Merriam-Webster dictionary, is defined as someone who is “prudently watchful and discreet in the face of danger or risk.”

falls too. As a consequence, even if the volatilities of consumption and business investment are not changing, total output is less volatile too. Interestingly, the effects of higher micro volatility and lower downpayments reinforce each other in explaining the reduced sensitivity of housing and hours to given aggregate shocks. This happens since lower downpayments have the effect of reducing, in the stationary equilibrium, the net worth of agents who are credit constrained and planning to buy a house; and because the dynamic effects of higher individual uncertainty are larger precisely when net worth is lower, since low income realizations are very costly in utility terms when individuals have, on average, a smaller buffer-stock of liquid assets.[44]

Figures 7 and 8 offer a graphical interpretation of these results. In figure 7, for each age, the two solid lines of each panel show how average debt, hours worked and housing positions in the lowest and the highest aggregate state. The top panel plots the calibration with high downpayment requirements and low idiosyncratic risk (the period 1952-1982). As the Figure shows, changes in the aggregate state using this calibration generate large differences in debt, housing and hours worked. The bottom panel plots the calibration with low required downpayments and high idiosyncratic risk (the period 1984-2008). Here, changes in the aggregate state generate smaller differences in debt, housing, and hours worked, thus illustrating how these variables become less volatile and less procyclical.

Figure 8 conveys the same information by plotting impulse responses for the key model variables when technology switches from its average value to a higher value (which corresponds roughly to 1 percent) in period 0. The responses are larger in the earlier period. On impact, residential investment falls before rising strongly in period 1. This result is well known in the household production literature (see, for instance, Greenwood and Hercowitz 1991 and Fisher 2007). In models with housing and business capital, business capital is useful for producing more types of goods than household 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 home to build up business capital, and only later is household capital accumulated. The important aspect to note is that higher idiosyncratic risk and lower downpayment requirements dampen the incentive to adjust housing capital, so that housing investment becomes less volatile.

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 agents' willingness to change housing consumption too much in presence of adjustment costs, while lower downpayment requirements allow households not to switch to smaller houses or to renting in bad times, since looser borrowing constraints provide a larger buffer against bad income shocks.

[44] Higher uncertainty in itself reduces the willingness to borrow, whereas lower downpayment lead to an increase in debt. In our baseline calibration, the second effect dominates – as shown in table 4, the ratio of debt to GDP rises from 0.29 to 0.35 when both changes are present. As a consequence, in the late period individuals are more cautious, even if they hold more debt.

6. Robustness Analysis

6.1. Four Alternative Versions of the Model

No Housing Adjustment Costs. A model without housing adjustment costs delivers a volatility of housing investment which is higher than in the baseline model. The standard deviation of residential investment, which is 6.69 percent in the baseline case, rises to 10.4 percent in absence of housing adjustment costs. Consumption volatility and fixed investment volatility are unchanged. Because houses are less risky, homeownership rises, from 64 to 68 percent. However, aggregate volatility is lower: because housing and nonhousing capital become closer substitutes as means of saving, the higher volatility of housing investment is more than offset by the reduced covariance between housing and nonhousing investment. In turn, this effect reduces the volatility of total output: the standard deviation falls from 2.09 percent to 2.02 percent.

Everybody rents. A version of the model in which everybody rents (we eliminate the utility cost of renting, so that $\theta = 1$ and renting is always preferred to owning) behaves similarly to a model without housing adjustment costs. The main difference with the model is that in this version there is no debt in equilibrium. In this model, the standard deviation of residential investment is 8.9 percent, while it is 6.69 percent 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 the required 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, 3.9 percent. Aggregate volatility is also lower: the standard deviation of GDP falls from 2.09 to 2.03 percent.

Homogeneous Discount Factor. In order to study the properties of the model with a homogeneous discount factor, we recalibrate our model so that it has the same homeownership rate and the same interest rate as in our baseline model. We change the discount factor (previously 0.925 for two-thirds of the population, and 0.995 for one-third) and the relative utility from renting (previously at 0.832). The values that achieve the same interest rate and homeownership rate of the baseline model are $\beta = 0.978$ and $\theta = 0.945$. At these parameter values, the volatility of housing investment and output are slightly higher than in the baseline calibration, but the correlation of housing investment with output and hours becomes smaller: this result occurs because fewer people are close to the borrowing limit and in need to increase hours worked to finance the downpayment in good times. In addition, with a single discount factor, very few people hold debt in equilibrium, and the distribution of wealth is more egalitarian than in the data: the Gini coefficient for wealth is 0.52, lower than in the data and in the benchmark model. Also, the correlation between debt and economic activity (which is about 0.8 in the data and the baseline model) drops to 0.18 only.

6.2. Some Final Thoughts

Relationship with the Literature on Lumpy Investment. Thomas (2002) has argued that lumpiness of fixed investment at the level of a single production unit (like a firm or a plant) bears no implications for the behavior of aggregate quantities in an otherwise standard equilibrium business cycle model. Her main argument rests on the representative household's desire to smooth its consumption profile over 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. It is not hard to see why: there is no representative household in our model. Adjustment costs imply smaller housing adjustment at the aggregate level, but larger housing adjustment (when it occurs) at the individual level.

Relationship with the Literature on Uncertainty and Durables Adjustment. Our result that higher uncertainty at the individual level reduces the volatility of aggregate housing investment echoes the results of papers in the literature that study how household durable purchases respond to changes in income uncertainty in (S, s) models resulting from transaction costs.[45] Eberly (1994), using data from the Survey of Consumer Finances, considers households' 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, Hurst and Leahy (2000) find a similar result using data on automobiles 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 – something that in our model is more likely to occur at high levels of m_h – 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.

Relationship with the Literature on Monetary Policy and Housing Investment. A variety of papers, policymakers and commentators have argued that it is nonsensical to think about fluctuations in housing investment while abstracting from interest rate shocks. According to this line of thought, fluctuations in the interest rate are the exogenous driving force behind business cycles, and housing dances to the tune of monetary policy. There is, however, little evidence for this claim in formal studies of the effect that monetary policy has on various components of aggregate demand. Bernanke and Gertler (1995), for instance, show that the initial percent decline in residential investment following an adverse interest rate shock is only marginally larger than the decline of business investment. Since the unconditional volatility of residential investment is about 1.5 larger than that of business investment, one can infer that the contribution of monetary policy shocks to fluctuations

[45] See Bertola, Guiso and Pistaferri (2005) for a full characterization of the role played by uncertainty in the adjustment of consumer durables.

in residential investment is not larger than the contribution of monetary policy to fluctuations in other types of investment. Because monetary policy shocks typically account for a small fraction of the fluctuations in economic activity,[46] we find it plausible to abstract from monetary policy in the context of a model as stylized as ours.[47]

7. Conclusions

In this paper, we formulate and solve an equilibrium business cycle model with household debt and 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 and household debt 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 stylized model, a result that is mirrored in the data. Our model accounts for the procyclicality of housing investment and for a good part of its relative variability, 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.

Despite its complexity, the model precludes an examination of certain aspects of housing behavior that may be important for understanding business cycle fluctuations. We have already mentioned in the previous section the role of monetary policy. Another limitation is that our assumptions about the technology, embedded in the economy's resource constraint, imply that the price of housing in units of consumption is constant. There are, however, three important reasons why we have not endogenized house prices in this paper. First and foremost, it is not obvious what extra insights this additional complication would buy: our main mechanisms rely on the interaction between risk, household debt, and housing demand, and movements in house prices should not dramatically affect these mechanisms. Second, 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 other shocks in addition to productivity shocks – like preference or sectoral shocks –, since we know from existing studies that technology shocks alone cannot quantitavely explain observed movements

[46] This fraction is typically 10 to 15 percent, measured through a variance decomposition of the forecast error of GDP explained by monetary shocks in vector autoregressions. See Christiano, Eichenbaum and Evans (1999).

[47] An intriguing possibility is that other episodes of financial innovation might have affected the transmission mechanism of monetary policy. Mertens (2008) shows how interest rate ceilings — such as those in place in the United States during the regulation Q period — can imply a larger effect of contractionary monetary policy on aggregate variables, so that their removal after the mid-1980s could have contributed to reduced business cycle volatility.

in house prices:[48] all of this would considerably increase computational costs.[49] Third, although movements in house prices are economically important, cyclical fluctuations in the price of housing are much smaller than the corresponding fluctuations in its quantity, which are, after all, the focus of our paper: for example, over the period 1970-2008, the standard deviation of year-on-year growth in real housing investment is 14 percent, while the corresponding number for real house prices is 3.7 percent.[50] For this reason, while we regard these issues as important,[51] we leave them for future work.

[48] See for instance Davis and Heathcote (2005) for a discussion.

[49] It takes about one day to solve our model on a Unix workstation.

[50] We use the Freddie Mac's Conventional Mortgage Home Price Index (adjusted for inflation), which starts in 1970. Of course, one might argue that volatility in house prices – especially a decline in house price – has more serious macroeconomic consequences than volatility of residential investment, but this issue is beyond the scope of this paper.

[51] The recent papers by Kiyotaki, Michaelides and Nikolov (2007), Favilukis, Ludvigson and Van Nieuwerburgh (2009), and Ríos-Rull and Sánchez-Marcos (2008) are important steps in this direction.

References

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 & Economic Statistics* 17(3), 373–81

Bernanke, Ben S, and Mark Gertler (1995) 'Inside the black box: The credit channel of monetary policy transmission.' *Journal of Economic Perspectives* 9(4), 27–48

Bertola, Giuseppe, Luigi Guiso, and Luigi Pistaferri (2005) 'Uncertainty and consumer durables adjustment.' *Review of Economic Studies* 72(4), 973–1007

Bottazzi, Renata, Hamish Low, and Matthew Wakefield (2007) 'Why do home owners work longer hours?' IFS Working Papers W07/10, Institute for Fiscal Studies, July

Budría, Santiago, Javier Diaz-Gimenez, Vincenzo Quadrini, and Jose-Victor Rios-Rull (2002) 'Updated facts on the u.s. distributions of earnings, income, and wealth.' *Quarterly Review* Summer, 2–35

Cagetti, Marco (2003) 'Wealth accumulation over the life cycle and precautionary savings.' *Journal of Business & Economic Statistics* 21(3), 339–53

Campbell, Jeffrey R., and Zvi Hercowitz (2005) 'The role of collateralized household debt in macroeconomic stabilization.' NBER Working Papers 11330, National Bureau of Economic Research, Inc, May

Carroll, Chris, and Wendy Dunn (1997) 'Unemployment expectations, jumping (s,s) triggers, and household balance sheets.' In 'NBER Macroeconomics Annual 1997, Volume 12' (MIT Press)

Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans (1999) 'Monetary policy shocks: What have we learned and to what end?' In *Handbook of Macroeconomics,* ed. J. B. Taylor and M. Woodford (Elsevier)

Davis, Morris A., and Jonathan Heathcote (2005) 'Housing and the business cycle.' *International Economic Review* 46(3), 751–784

Domeij, David, and Martin Floden (2006) 'The labor-supply elasticity and borrowing constraints: Why estimates are biased.' *Review of Economic Dynamics* 9(2), 242–262

Dynan, Karen E., Douglas W. Elmendorf, and Daniel E. Sichel (2007) 'The evolution of household income volatility.' Finance and Economics Discussion Series 2007-61, Board of Governors of the Federal Reserve System (U.S.)

Eberly, Janice C (1994) 'Adjustment of consumers' durables stocks: Evidence from automobile purchases.' *Journal of Political Economy* 102(3), 403–36

Favilukis, Jack, Sydney Ludvigson, and Stijn Van Nieuwerburgh (2009) 'Macroeconomic implications of housing wealth, housing finance, and limited risk-sharing in general equilibrium.' Unpublished

Fernandez-Villaverde, Jesus, and Dirk Krueger (2004) 'Consumption and saving over the life cycle: How important are consumer durables?' 2004 Meeting Papers 357b, Society for Economic Dynamics

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

Fisher, Jonas D. M., and Martin Gervais (2007) 'First-time home buyers and residential investment volatility.' Working Paper Series WP-07-15, Federal Reserve Bank of Chicago

Foote, Christopher, Erik Hurst, and John Leahy (2000) 'Testing the (s, s) model.' *American Economic Review* 90(2), 116–119

Fortin, Nicole M (1995) 'Allocation inflexibilities, female labor supply, and housing assets accumulation: Are women working to pay the mortgage?' *Journal of Labor Economics* 13(3), 524–57

Gerardi, Kristopher, Harvey Rosen, and Paul S. Willen (forthcoming) 'The impact of deregulation and financial innovation on consumers: The case of the mortgage market.' *Journal of Finance*

Gervais, Martin (2002) 'Housing taxation and capital accumulation.' *Journal of Monetary Economics* 49(7), 1461–1489

Gomme, Paul, Richard Rogerson, Peter Rupert, and Randall Wright (2004) 'The business cycle and the life cycle.' In 'NBER Macroeconomics Annual 2004, Volume 19' (MIT Press)

Greenwood, Jeremy, and Zvi Hercowitz (1991) 'The allocation of capital and time over the business cycle.' *Journal of Political Economy* 99(6), 1188–214

Gruber, Joseph, and Robert Martin (2003) 'Precautionary savings and the wealth distribution with illiquid durables.' International Finance Discussion Papers 773, Board of Governors of the Federal Reserve System (U.S.)

Haider, Steven J (2001) 'Earnings instability and earnings inequality of males in the united states: 1967-1991.' *Journal of Labor Economics* 19(4), 799–836

Hansen, Gary D (1993) 'The cyclical and secular behaviour of the labour input: Comparing efficiency units and hours worked.' *Journal of Applied Econometrics* 8(1), 71–80

Hansen, Kristin A. (1998) 'Seasonality of moves and duration of residence.' Current Population Reports, Household Economic Studies, P70-66 (October), U.S. Census Bureau

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante (2008) 'The macroeconomic implications of rising wage inequality in the united states.' NBER Working Papers 14052, National Bureau of Economic Research, Inc, June

Heaton, John, and Deborah J Lucas (1996) 'Evaluating the effects of incomplete markets on risk sharing and asset pricing.' *Journal of Political Economy* 104(3), 443–87

Hendricks, Lutz (2007) 'How important is discount rate heterogeneity for wealth inequality?' *Journal of Economic Dynamics and Control* 31(9), 3042–3068

Kennickell, Arthur, and Janice Shack-Marquez (1992) 'Changes in family finances from 1983 to 1989: evidence from the survey of consumer finances.' *Federal Reserve Bulletin* (Jan), 1–18

King, Robert G., and Sergio T. Rebelo (1999) 'Resuscitating real business cycles.' In *Handbook of Macroeconomics,* ed. J. B. Taylor and M. Woodford, 1 ed., vol. 1 (Elsevier) chapter 14, pp. 927–1007

Kiyotaki, Nobuhiro, Alexander Michaelides, and Kalin Nikolov (2007) 'Winners and losers in housing markets.' CDMA Conference Paper Series 0705, Centre for Dynamic Macroeconomic Analysis, November

Krueger, Dirk, and Fabrizio Perri (2006) 'Does income inequality lead to consumption inequality? evidence and theory.' *Review of Economic Studies* 73(1), 163–193

Krusell, Per, and Anthony A. Smith (1998) 'Income and wealth heterogeneity in the macroeconomy.' *Journal of Political Economy* 106(5), 867–896

Mertens, Karel (2008) 'Deposit rate ceilings and monetary transmission in the us.' *Journal of Monetary Economics* 55(7), 1290–1302

Moffitt, Robert, and Peter Gottschalk (2008) 'Trends in the transitory variance of male earnings in the u.s., 1970-2004.' Boston College Working Papers in Economics 697, Boston College Department of Economics, December

Nakajima, Makoto (2005) 'Rising earnings instability and prices of housing and financial assets.' 2005 Meeting Papers 211, Society for Economic Dynamics

Ortalo-Magné, François, and Sven Rady (2006) 'Housing market dynamics: On the contribution of income shocks and credit constraints.' *Review of Economic Studies* 73(2), 459–485

Poterba, James M. (1992) 'Taxation and housing markets.' In 'Canada-U.S. Tax Comparisons' (University of Chicago Press)

Queisser, Monika, and Edward Whitehouse (2005) 'Pensions at a glance: public policies across oecd countries.' MPRA Paper 10907, University Library of Munich, Germany, May

Ríos-Rull, José-Víctor (1996) 'Life-cycle economies and aggregate fluctuations.' *Review of Economic Studies* 63(3), 465–89

Ríos-Rull, José-Víctor, and Virginia Sánchez-Marcos (2008) 'An aggregate economy with different size houses.' *Journal of the European Economic Association* 6(2-3), 705–714

Rosen, Harvey S. (1985) 'Housing subsidies: Effects on housing decisions, efficiency, and equity.' In *Handbook of Public Economics,* ed. A. J. Auerbach and M. Feldstein, vol. Handbook of Public Economics (North Holland)

Scholz, John Karl, Ananth Seshadri, and Surachai Khitatrakun (2006) 'Are americans saving "optimally" for retirement?' *Journal of Political Economy* 114(4), 607–643

Silos, Pedro (2007) 'Housing, portfolio choice and the macroeconomy.' *Journal of Economic Dynamics and Control* 31(8), 2774–2801

Stock, James H., and Mark W. Watson (2002) 'Has the business cycle changed and why?' In 'NBER Macroeconomics Annual 2002, Volume 17' (The MIT Press)

Storesletten, Kjetil, Chris I. Telmer, and Amir Yaron (2004) 'Cyclical dynamics in idiosyncratic labor market risk.' *Journal of Political Economy* 112(3), 695–717

Tauchen, George (1986) 'Finite state markov-chain approximations to univariate and vector autoregressions.' *Economics Letters* 20(2), 177–181

Thomas, Julia K. (2002) 'Is lumpy investment relevant for the business cycle?' *Journal of Political Economy* 110(3), 508–534

Van Nieuwerburgh, Stijn, and Pierre-Olivier Weill (2006) 'Why has house price dispersion gone up?' NBER Working Papers 12538, National Bureau of Economic Research, Inc, September

Venti, Steven F., and David A. Wise (2001) 'Choice, chance, and wealth dispersion at retirement.' In 'Aging Issues in the United States and Japan' (University of Chicago Press)

Ziliak, James P., and Thomas J. Kniesner (1999) 'Estimating life cycle labor supply tax effects.' *Journal of Political Economy* 107(2), 326–359

Appendix A: Computational Details

We 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 (with the exception of the policy functions for housing, that are defined only on points on the grid). The algorithm consists of the following steps:

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

   The number of grid points was chosen as follows: 7 points for the aggregate shock, 3 values for the idiosyncratic shock, 15 points for the housing stock, and 350 points for the financial asset.[52] For aggregate capital, we choose a grid of 20 equally spaced points in the initial range $[0.8K^*, 1.2K^*]$, where K^* denotes the average value of this variable in the simulations. The range is then updated at each iteration consistently with the simulated K, 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 functions that approximate the laws of motion of capital and labor:

$$\begin{aligned} K_t &= \omega_0^A + \omega_1^A K_{t-1}, \\ L_t &= \omega_2^A + \omega_3^A K_{t-1}. \end{aligned}$$

   Because factor prices (wages and interest rates) only depend on aggregate capital and labor in equilibrium, this approach is equivalent to assuming that individuals forecast these factor prices using a function of K_{t-1} for wach value of the aggregate state A.

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.[53] 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:[54]

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

   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:

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

[52] 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.

[53] 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).

[54] We prevent individuals from choosing negative hours.

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:

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

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. 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 90,000 individuals for 5,000 periods, discarding the first 500 periods.[55] Compute the aggregate variables K and L at each t.

5. Run a regression of the simulated aggregate capital and the simulated aggregate labor on lagged aggregate capital, retrieving the new coefficients $\left\{\omega_i^A\right\}$ for the laws of motion for K and L. We repeat steps 3 and 4 until convergence over the coefficients of the regressions. We measure goodness of fit using the R^2 of the regressions: they are always equal to 0.997 or higher at convergence for K and around 0.95 for L; the corresponding wage rate and interest rate functions are also very accurate: the R^2 of the regression of the wage rate on aggregate K is 0.999, the R^2 of the regression of the interest rate on aggregate K is 0.992.

[55] 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 Volatility

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

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

We want to pick values for ρ_Z and σ_Z that are in line with evidence. Below is a selected survey of studies that have attempted to estimate these parameters in setups similar to ours.[56]

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. After controlling for fixed effects (η), for the period 1969 to 1984 (see their Table A.3), they estimate using PSID data the following values: $\rho_Z = 0.53$ and $\sigma_Z = 0.296$.

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 their Appendix A) for married, no college, two-earners are $\rho_Z = 0.699$ and $\sigma_Z = 0.428$.

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 with a regime-switching conditional variance of persistent shocks. Setting aside their transitory shocks, their estimates of ρ_Z are around 0.95, and their estimated standard deviation of the innovation to the persistent component is 0.138.[57] This translates in an unconditional standard deviation of log income equal to $\sigma_Z = 0.138/\left(1 - 0.95^2\right)^{1/2} = 0.44$.

The Change in Volatility

Several studies document the increase in the cross-sectional dispersion of earnings in the United States between the 1970s and the 1990s. 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.

1. Moffitt and Gottschalk (2008) study changes in the variance of permanent and transitory component of income volatility using data from the PSID from 1970 to 2004. They find that the non-permanent component (transitory) variance of earnings (for male workers) increased substantially in the 1980s and then remained at this new higher level through 2004. They report (see Figure 7 in their paper) that the variance of the transitory component rose from around 0.10 to 0.22 between the 1970s and the 1980s-1990s. This corresponds to a rise in the standard deviation from 0.32 to 0.47. Their estimate of the autocorrelation of the transitory shocks is 0.85.

[56] 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. We do not attempt to control for this in our calibration.

[57] We calculate this number using a weighted average of their reported standard deviation across regimes, which are 0.12 in expansions, and 0.21 in recessions, using the observation that expansions last on average four times longer than recessions in postwar U.S. data.

2. 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. They find that the variance of persistent shocks doubles during the 1975-1985 decade.

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_Z is 0.639 (Table 4 in his paper). He finds that the unconditional standard deviation of the instability component rises from around $0.23 - 0.24$ to about $0.35 - 0.37$ during the 1980s.

4. Krueger and Perri (2006) model log income as an ARMA process of the kind

$$\begin{aligned} y_t &= z_t + \varepsilon_t, \\ z_t &= \rho_Z z_{t-1} + \sigma_Z \left(1 - \rho_Z^2\right)^{1/2} \varepsilon_t^z, \\ \varepsilon_t &= \sigma_\varepsilon \varepsilon_t^e \end{aligned}$$

where ε_t^e and ε_t^z are $Normal\,(0,1)$. They allow the innovation variances σ_ε and σ_Z to vary by year. They find that the values of σ_Z and σ_ε are respectively 0.42 and 0.28 in 1980, and 0.52 and 0.36 in 2003. Given these numbers, the standard deviation of log income y_t rises by 0.13, from $\sqrt{0.42^2 + 0.28^2} = 0.50$ to $\sqrt{0.52^2 + 0.36^2} = 0.63$.

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. They report that the standard deviation of the innovations to income growth rose from 0.2 to 0.25, approximately.

From this brief review, we conclude that a plausible value for the persistence of the income shock is around 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 in the range of estimates reported by Moffitt and Gottschalk (2008).

Table 1. U.S. Economy. Cyclical Statistics.

	1952.I -1982.IV (Early Period)			1984.I -2008.IV (Late Period)		
	s.d. %	ratio[i]	corr. w/ GDP	s.d. %	ratio[i]	corr. w/ GDP
GDP[ii]	2.09	1.00	1.00	1.27	1.00	1.00
C[iii]	1.20	0.57	0.92	0.61	0.48	0.91
IH	8.24	3.94	0.84	3.78	2.97	0.72
IK	5.03	2.41	0.77	4.14	3.25	0.84
Debt	2.23	1.21	0.78	1.79	1.41	0.29

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*) Durables expenditures are assigned to IH. All series are in logs and detrended with HP-filter with smoothing parameter 1,600.

Table 2: Parameter Values for the Benchmark Model Economy

	Parameter	Value	Target/Source
Preferences			
Discount factor, patient agents	β_H	0.995	$R = 3.75\%$
Discount factor, impatient agents	β_L	0.925	-
Fraction of impatient agents	-	2/3	Gini wealth 0.76
Relative weight on leisure in utility	τ	1.5	-
Productive time	$\bar{l}$	2.5	Time worked 40%
Relative weight on housing in utility	j	0.15	$H/Y = 1.22$
Relative utility from renting vs. owning	θ	0.832	Homeownership 64%
Utility weights (family size)	λ_a	see text	-
Life, retirement			
Survival probabilities	Π_a	see text	U.S. decennial life tables
Retirement age	$\widetilde{T}$	65	-
Pension	P	0.4	40% average wage
Technology			
Capital share	α	0.33	$K/Y = 2.4$
Business capital depreciation rate	δ_K	0.1	$IK/Y = 0.24$
Housing depreciation rate	δ_H	0.05	$IH/Y = 0.06$
Autocorrelation of tech. shock	ρ_A	0.925	King and Rebelo (1999)
St.dev. (unconditional) of techn. shock	σ_A	0.0139	U.S. filtered GDP st.dev.
Housing transaction cost	ψ	4%	-
Minimum House Size	$\underline{h}$	2 × avg wage	U.S. Stat.Abst.(Tab.948)
Borrowing			
Max borrowing, fraction of lifetime wage	m_y	0.25	-
Maximum borrowing, fraction of house	m_h	0.75	-
Individual income process			
Autocorrelation of earnings shock	ρ_Z	0.90	see appendix B
St.dev. (unconditional) of earnings shock	σ_Z	0.30	see appendix B
Age-dependent earnings ability	Hansen (1993), quadratic interpolation		

Table 3: U.S. Economy and Model. Cyclical Statistics. Comparison for the Early Period.

	DATA: 1952.I -1982.IV			MODEL		
	s.d. %	ratio[i]	corr. w/ GDP	s.d.%	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	1.67	0.80	0.96
IH	8.24	3.94	0.84	6.69	3.20	0.60
IK	5.03	2.41	0.77	3.66	1.75	0.83
$Debt^{\text{iv}}$	2.23	1.06	0.78	9.77	4.68	0.82

Notes: (*i*) The ratio is the standard deviation of the variable scaled by that of GDP; (*ii*) C, IH and IK are logged, 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. The model moments are based on statistics from a simulation of 5,000 periods. The data are detrended using HP-filter with $\lambda = 1,600$.

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 $m_h = 0.85$ $\sigma_Z = 0.45$
s.d.%				
GDP	2.09	2.06	2.06	2.02
C	1.67	1.64	1.69	1.70
IH	6.69	6.41	6.30	5.34
IK	3.66	3.45	3.69	3.55
$Debt$	9.77	3.47	6.92	1.30
Corr$(Debt, GDP)$	0.82	0.66	0.78	0.35
% $Homeown$	64%	76%	62%	67%
Debt to GDP	0.29	0.46	0.26	0.35
$\sigma\left(\Delta C_{it}/C_i\right)$	0.63	0.76	0.62	0.67
Gini wealth	0.76	0.76	0.77	0.77
Gini labor income	0.42	0.42	0.49	0.49
Gini consumption	0.26	0.27	0.31	0.31

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), we increase the loan-to-value ratio from 0.75 to 0.85. In (3), we increase earnings volatility from 0.3 to 0.45. In (4), we increases both loan-to-value ratio and earnings volatility and calibrate the U.S. economy for the period 1984-2008.

Figure 1: Household Debt, Housing Investment and GDP (HP-filtered Variables).

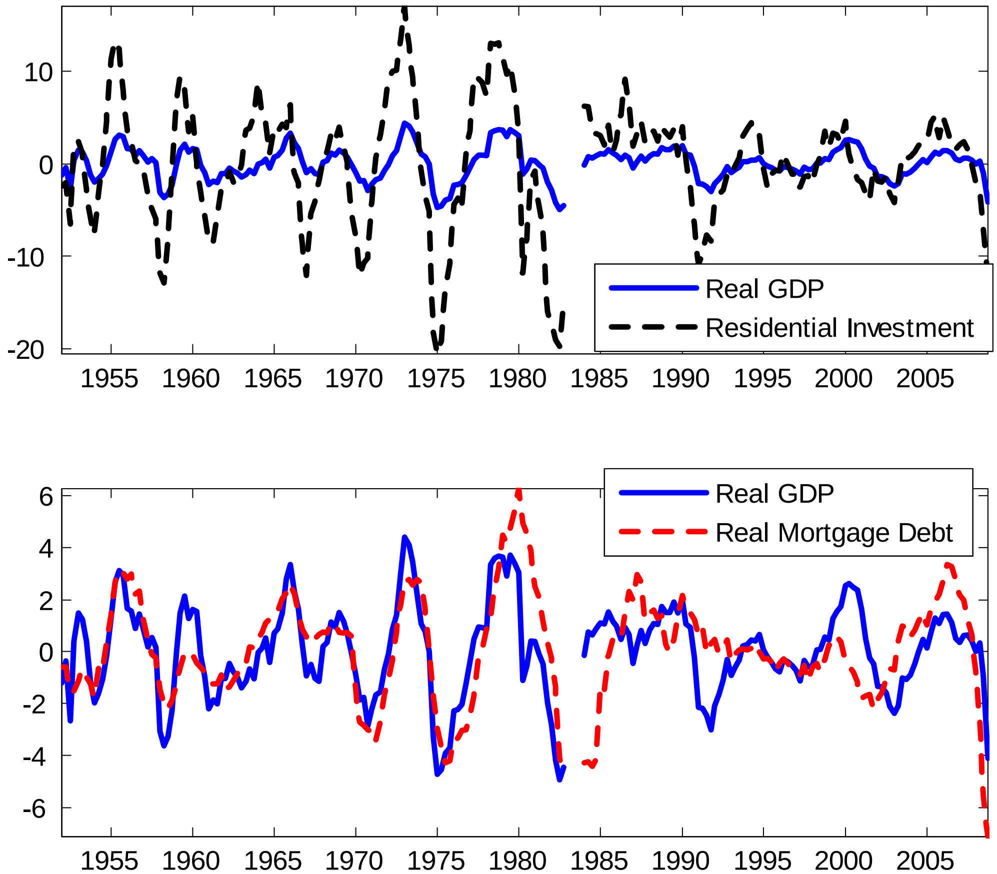

Note: Variables are expressed in percentage deviation from their trend.

Figure 2: Homeowner's Housing Investment Decision as a Function of Initial House Size and Liquid Assets.

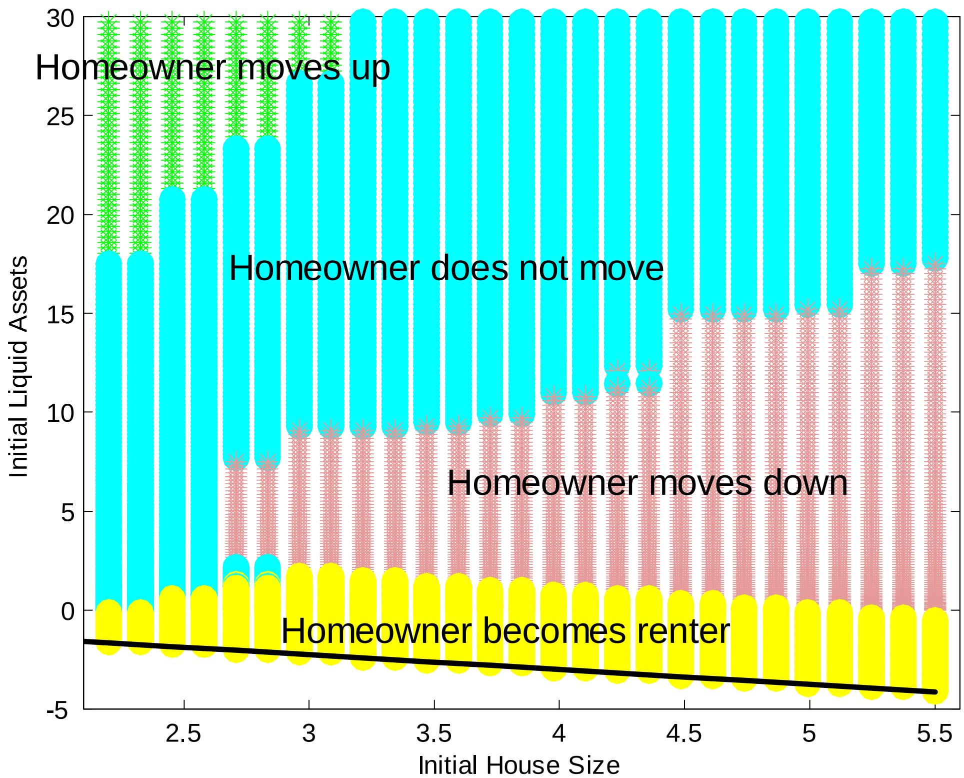

Note: The figure illustrates, for each combination of initial house and liquid assets, the homeowner's housing decision for next period. It is plotted for a patient agent who is 65 years old, when aggregate productivity and the average capital labor ratio are equal to their average value.

Figure 3: A Typical Life-cycle Profile.

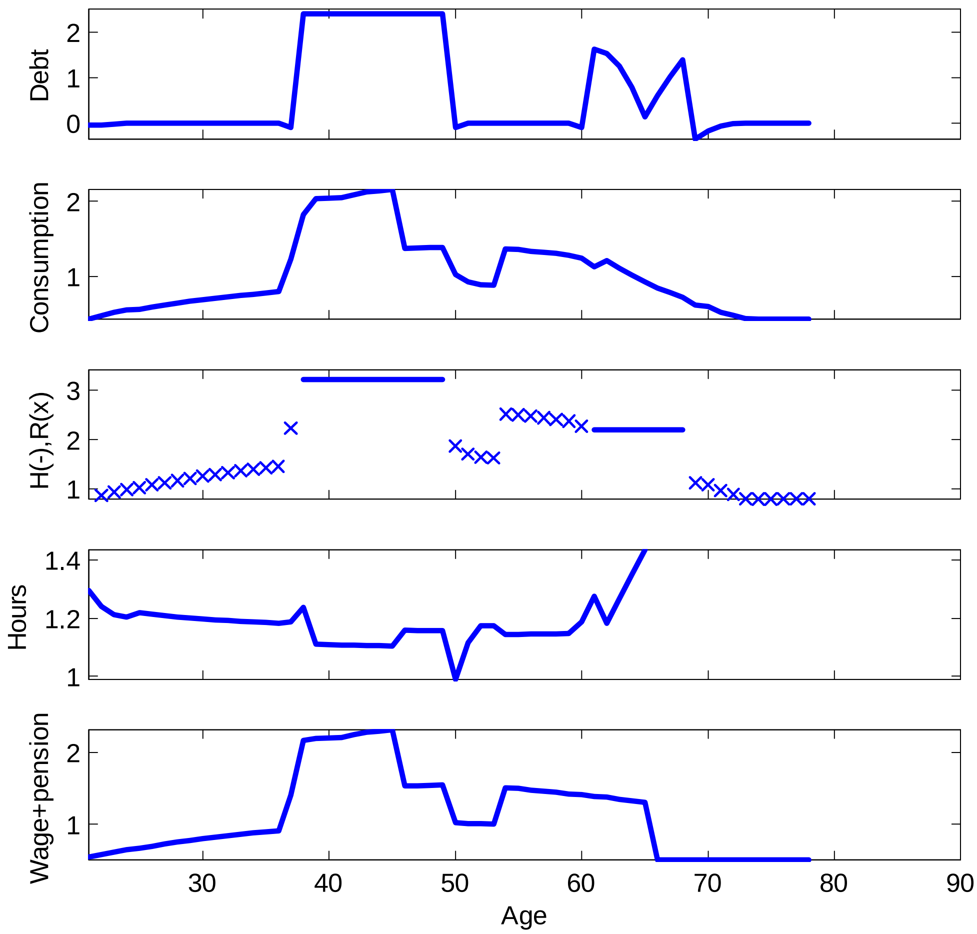

Note: This figure plots life-cycle choices of a randomly chosen impatient agent from birth (age 21) to death (age 78). In panel 3, the "x" symbol denotes the amount rented when the individual is renting, whereas the solid line denotes the amount owned when the individual owns a house.

Figure 4: Comparison between Model (Baseline Calibration) and Data (1983 Survey of Consumer Finances).

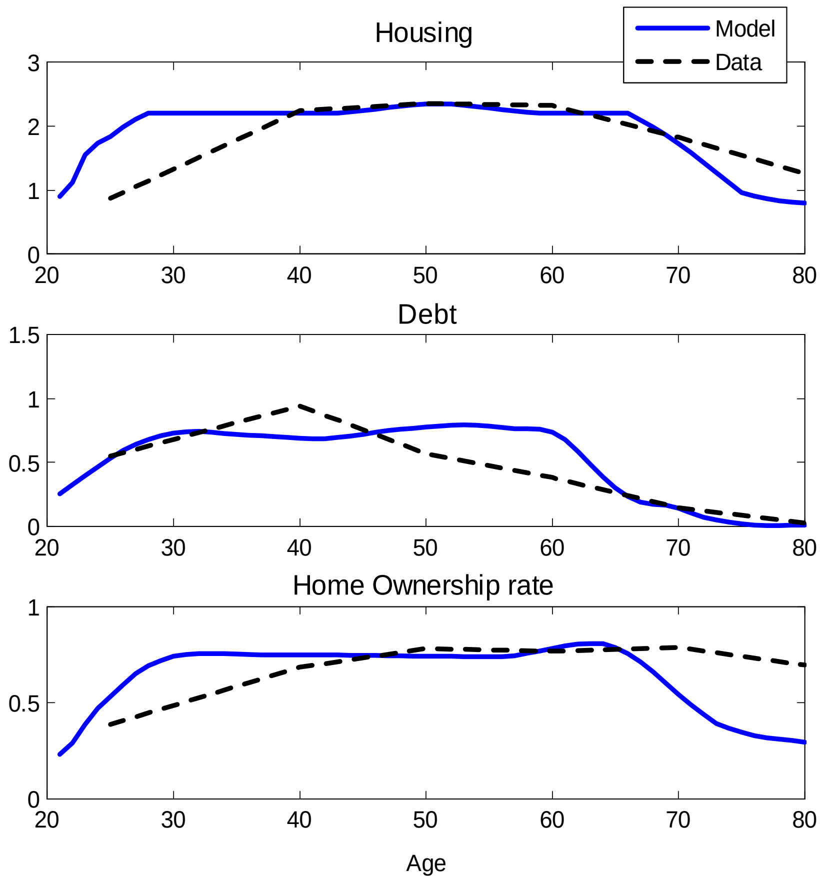

Note: The data come from the summary statistics of the 1983 Survey of Consumer Finances, as reported in Kennickell and Shack-Marquez (1992). For each age, the model variable is the product of the fraction of households in that age holding housing or debt, times the median holding of housing or debt. The data variable is constructed in the same way.

Figure 5: Lorenz Curves for Wealth, Owner-occupied Housing and Consumption.

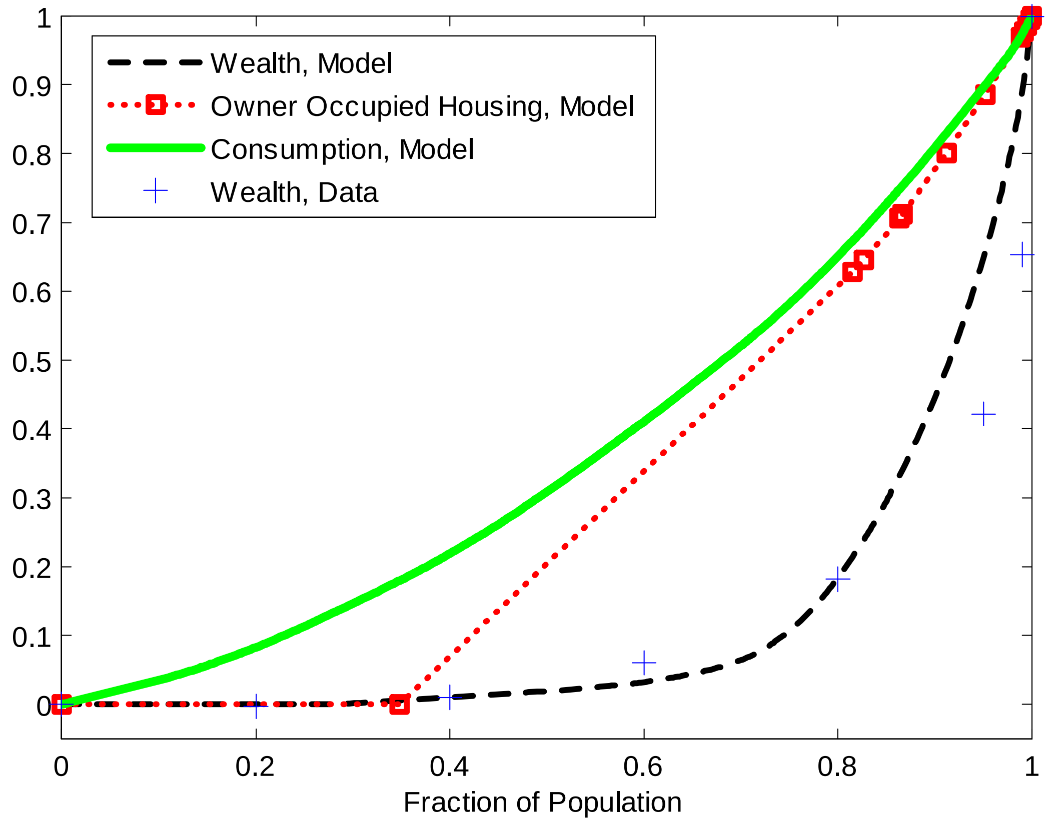

Note: The + sign refers to the data (source: Budría et al., 2002).

Figure 6: Equilibrium Hours Worked for Renters, Borrowers and Savers in Response to Wage Changes.

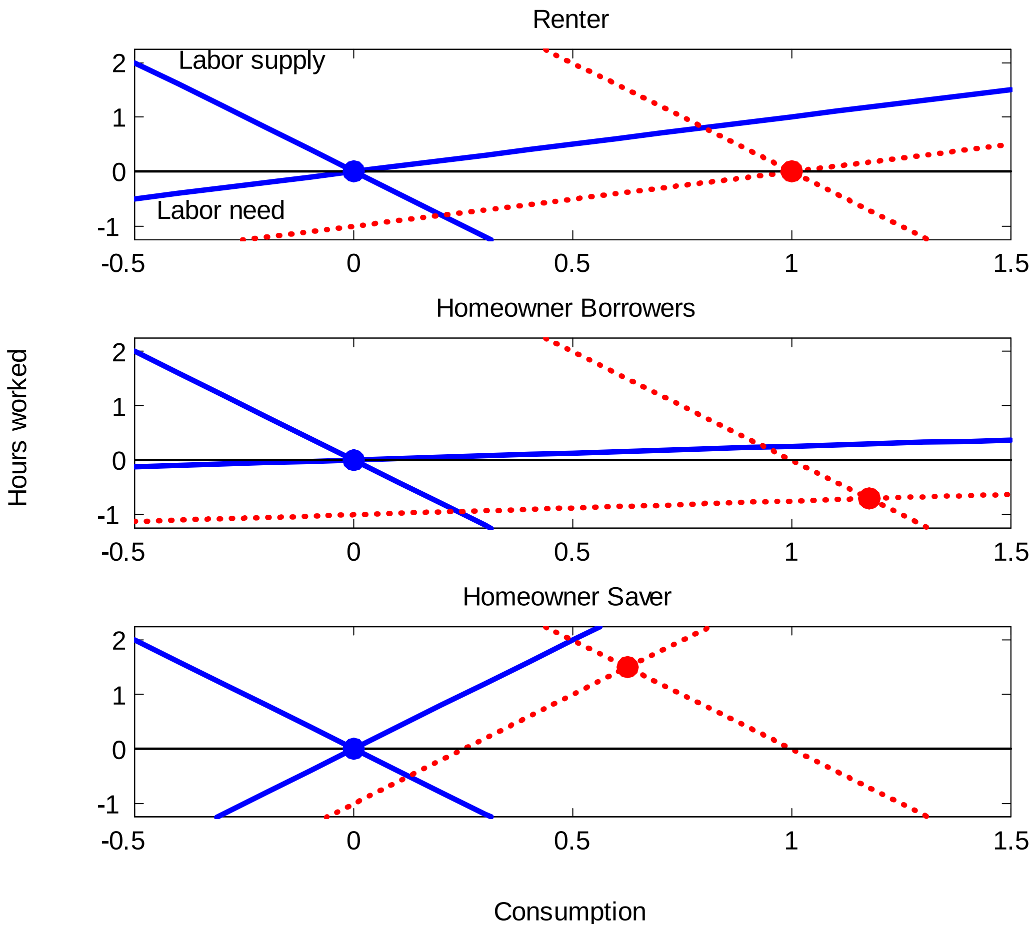

Note: The horizontal and vertical axis plot respectively percentage deviations of consumption and hours from their steady state values. The downward sloping line plots is the labor supply curve as a function of consumption (the negative slope reflects the negative wealth effect on labor supply from higher 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 7: Comparison between Early and Late Period: Debt, Hours and Housing by Age.

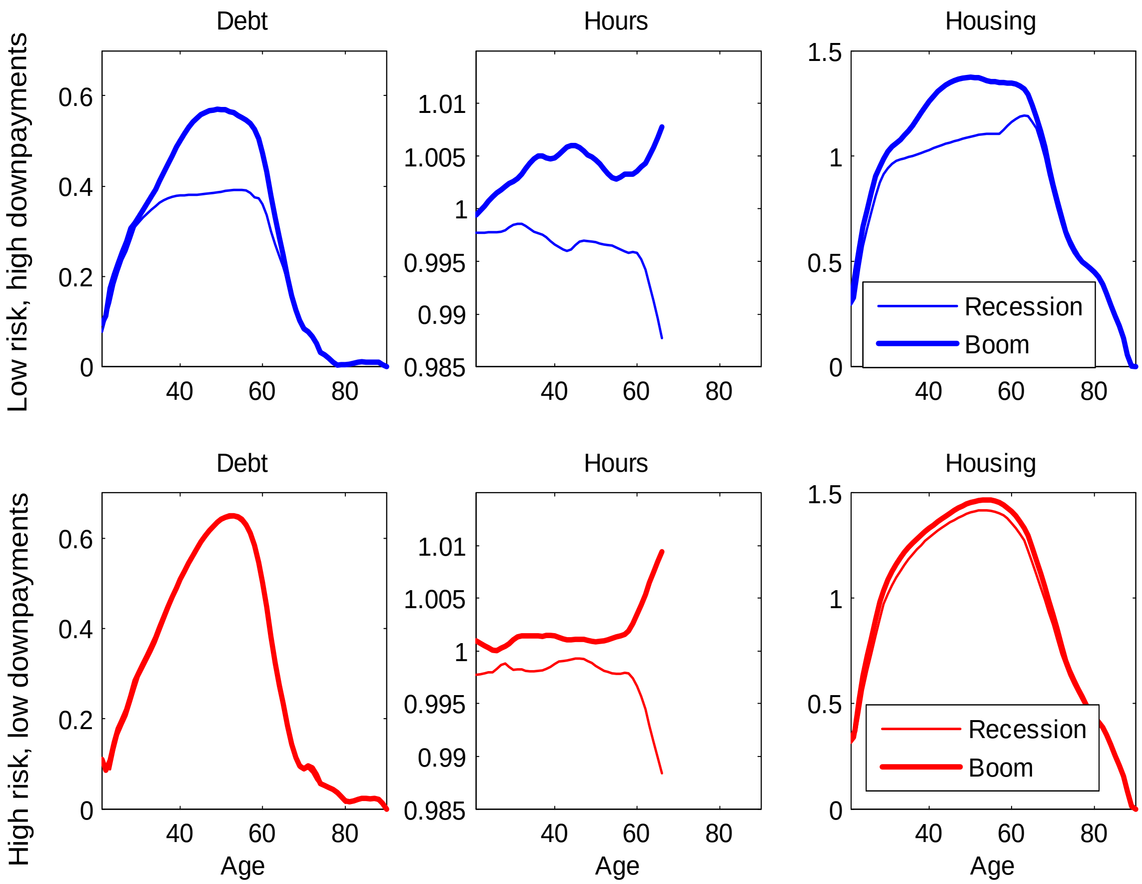

Note: The top panel plots model variables in the baseline calibration (low individual risk and high downpayment requirements), where housing, debt and hours worked are relatively more volatile (the difference between a boom and a recession is larger). The bottom panel plots the calibration with high individual risk and low downpayment requirements.

The thin/thick line shows the reading of each variable by age when the economy in the lowest/highest aggregate state. Consumption, Housing and Debt are expressed as a ratio of average GDP. Hours are normalized in each age by their age average.

Figure 8: Impulse Responses to a Technology Shock: Comparison between the Early Period and Late Period.

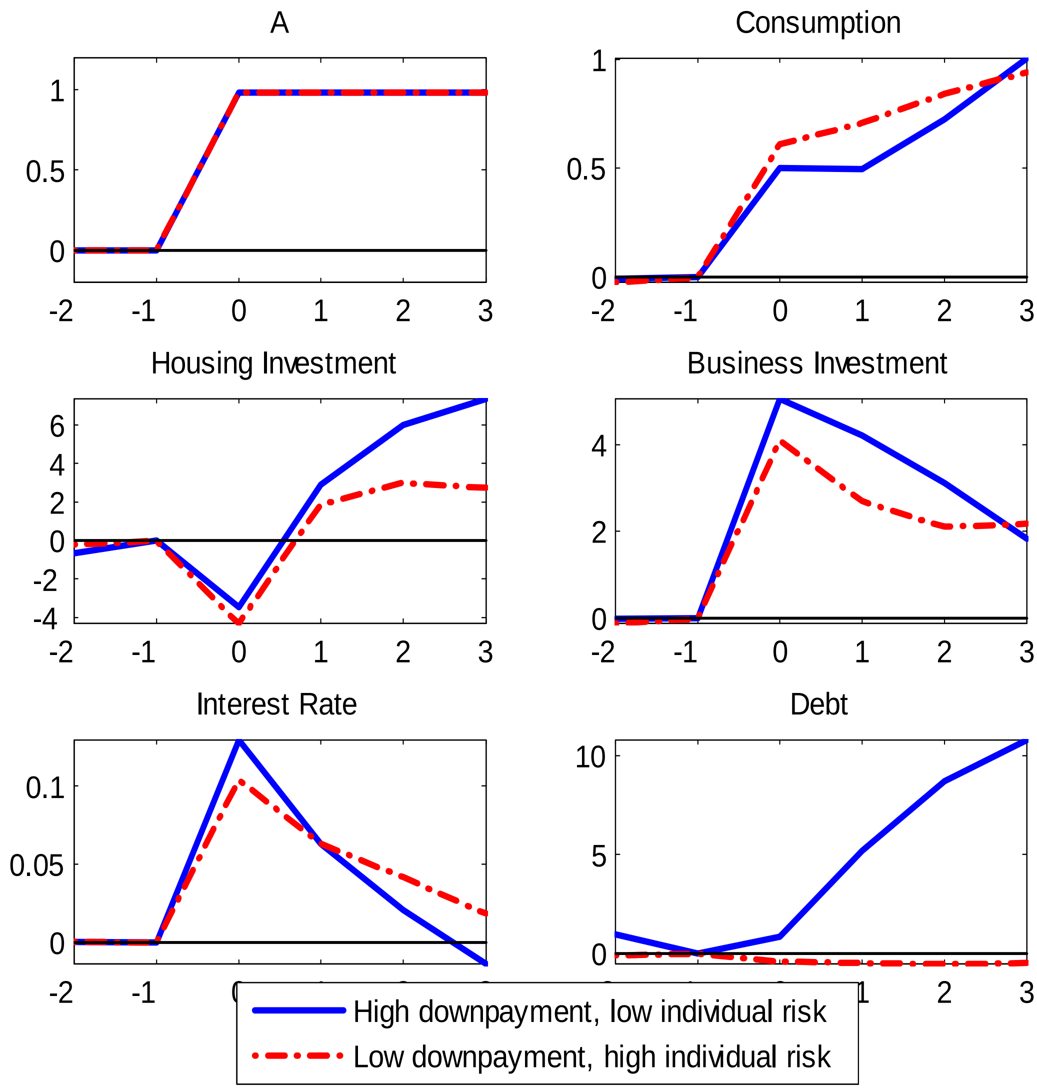

Note: Model dynamics following an exogenous switch in aggregate productivity A (in period zero) from the median state to next higher value (a 0.97 percent increase) that lasts at least four periods. Each variable is displayed in percentage deviation from the unshocked path.