

INFLATION AND THE PRICE OF REAL ASSETS*

Matteo Leombroni Monika Piazzesi Ciaran Rogers Martin Schneider

January 2020

Abstract

In the 1970s, U.S. asset markets witnessed (i) a 25% dip in the ratio of aggregate household wealth relative to GDP and (ii) negative comovement of house and stock prices that drove a 20% portfolio shift out of equity into real estate. This study uses an overlapping generations model with uninsurable nominal risk to quantify the role of structural change in these events. We attribute the dip in wealth to the entry of baby boomers into asset markets, and to the erosion of bond portfolios by surprise inflation, both of which lowered the overall propensity to save. We also show that the Great Inflation led to a portfolio shift by making housing more attractive than equity. Disagreement about inflation across age groups matters for the size of tax effects, the volume of nominal credit, and the price of housing as collateral.

*Department of Economics, Stanford University, leombm@stanford.edu, piazzesi@stanford.edu, ciaran@stanford, schneidr@stanford. For comments and suggestions, we thank Joao Cocco, Jesus Fernandez-Villaverde, John Heaton, Susan Hume McIntosh, Larry Jones, Patrick Kehoe, Per Krusell, Ricardo Lagos, Ellen McGrattan, Toby Moskowitz, Neng Wang, and many seminar and conference participants.

I Introduction

The 1970s brought dramatic changes in the size and composition of US household sector wealth. Figure 1 plots key household sector positions over the postwar period. During the 1970s, household net worth as a fraction of GDP fell by 25%, before recovering again to its late 1960s value. Figure 2 zooms in on this era and shows portfolio *weights* in Panel (a), in particular a 20pp shift away from equity and into real estate during. This portfolio adjustment was largely driven by negative comovement of asset prices – house prices rose while equity prices fell, as shown in Panel (b). Compared to the big swings in the major real asset positions, households’ net position in nominal credit instruments was relatively stable. As documented below, the stability of net positions masks substantial increases in gross borrowing and lending within the household sector.

This paper studies the turbulence of the 1970s using an asset pricing model with heterogeneous agents and incomplete markets. We consider overlapping generations of households who invest in equity and housing and who borrow and lend in nominal credit instruments. We fit the model to micro data on US households’ income and portfolio positions and then perform counterfactuals to understand the 1970s. We attribute the wide swings in asset prices – and thereby household wealth – to two events unique to that decade: the entry of the young baby boomers into asset markets and the Great Inflation that altered the distribution of real returns on credit for nominal borrowers and lenders. We show how those events contribute to a lower average savings rate as well as a portfolio shift towards housing for the aggregate household sector. At the same time, the outstanding quantity of houses, equity and government debt declined only slightly; shifts in household asset demand thus resulted in large price adjustments.

Individual asset demand in our model comes from a life cycle consumption savings problem with uninsurable income risk and portfolio choice between housing, equity and nominal bonds. Households can borrow only against housing, not against equity or human wealth (that is, permanent income). This feature helps to match that (i) many households have a lot debt, but positive net worth and (ii) younger and poorer households who have low financial wealth relative to permanent income save less out of financial wealth. We further work with a standard income process: most income risk is idiosyncratic so human wealth is not very exposed to aggregate shocks moving house and stock prices. This feature helps match observed portfolio choice: younger households build riskier financial portfolios, in particular through leveraged investment in housing. This is optimal because they have relatively more human wealth that is not exposed to asset price movements.

Our asset pricing results reflect how demographics and the Great Inflation altered both savings and portfolio choice. First, savings declined because there were relatively more young households, and because surprise inflation eroded the financial wealth of nominal lenders. Second, beliefs about future inflation induced a portfolio shift towards housing. A prominent reason is that nominal rigidities in the tax code favor housing over equity in times of high expected inflation. Moreover, survey data show that households disagreed about expected inflation: young borrowers expected higher inflation – and hence lower real interest rates – than old lenders. The resulting gains from trade in the nominal credit

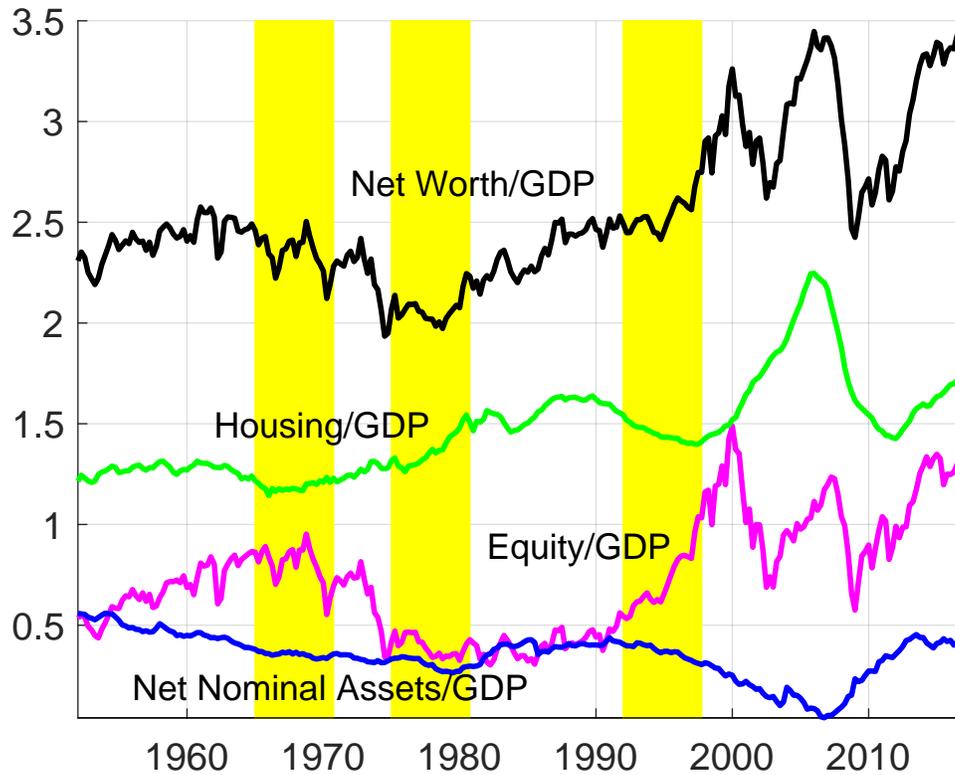
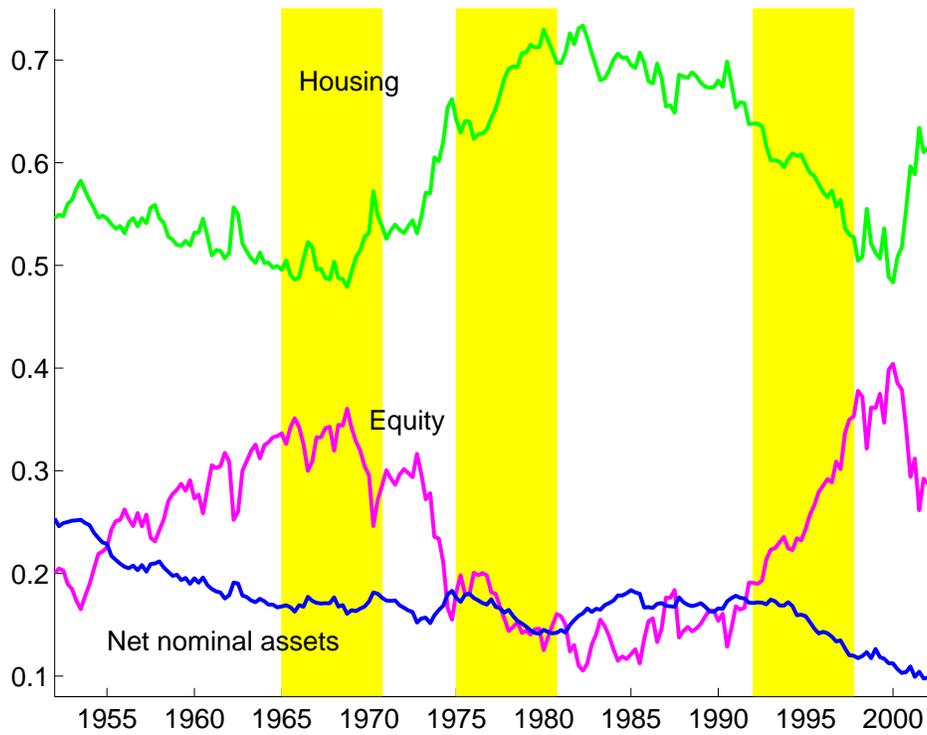


Figure 1: Aggregate wealth components (market value) divided by GDP, Financial Accounts of the United States & own computations, 1952:1-2016:4.

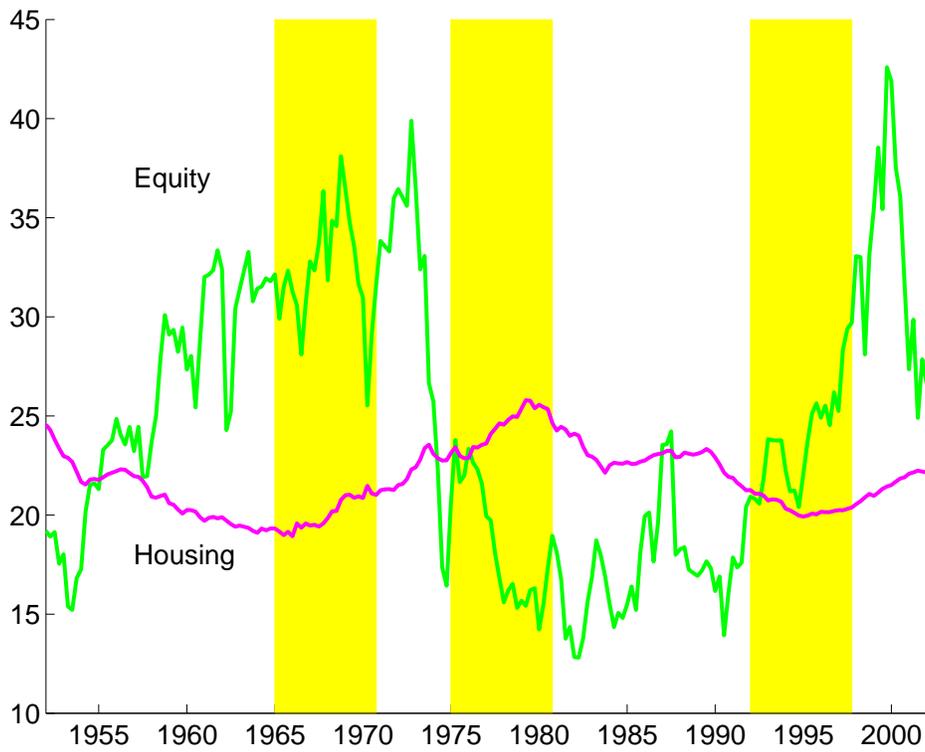
market increased credit and pushed up the price of collateral. At the same time, inflation uncertainty provided an important counterforce that dampened the response of credit.

Our empirical strategy uses micro data on household characteristics to directly parameterize household sector asset demand. Our model describes three 6-year trading periods, highlighted in yellow in Figures 1-2. For each trading period, we estimate the distribution of income and initial asset endowments across households from the Survey of Consumer Finances (SCF). We parameterize beliefs about future prices and income using a mix of historical data and survey forecasts. We can then compute optimal household asset demand functions given prices, endowments and expectations. Finally, we find equilibrium prices for the trading period to equate aggregate household asset demand to the (observed) quantity of assets supplied to households by other sectors, measured using the Financial Accounts of the United States (FAUS).

To establish that the model is a useful tool for thinking about asset pricing with heterogeneous agents, we check its ability to match not only prices but a rich cross section of household savings and portfolio choices. We select preference parameters so the model closely replicates the aggregate wealth-to-GDP ratio, household sector portfolio weights, as well as the nominal interest rate in 1995. It then also matches price dividend ratios for both equity and real estate. The success of the model is that it



(a) Aggregate Portfolio Shares



(b) Price-dividend ratio and price-rent ratio

Figure 2: Financial Accounts of the United States & own computations, 1952:1-2003:4

captures a number of (untargeted) stylized facts about the cross-section of households in the 1995 SCF. In particular, it generates hump-shaped cohort market shares in wealth, real estate and equity, as well as net nominal positions that increase – and real estate shares that decrease – with age and net worth. The main mechanism behind these facts is that agents who expect more future non-asset income are willing to build more risky portfolios.

Counterfactuals that vary endowments, supply or expectations then serve to show the quantitative importance of each element for price fluctuations. In particular, we use the model to examine the 1970s. We first show that it produces a drop in aggregate wealth between 1968 and 1978 for a wide range of expectation scenarios. It attributes the dip in the wealth-GDP ratio to two effects. First, the entry of baby boomers into asset markets lowered the average savings rate. Second, capital losses from realized inflation lowered wealth and hence savings, especially for older households. At the same time, lower savings were not counteracted by a large increase in interest rates, because the outside supply of bonds to the household sector also fell – the main force here is a reduction of government debt. Here, measurement of asset supply is crucial for the quantitative results.

Our quantitative analysis implies that both inflation expectations and inflation uncertainty were relevant for asset prices and household positions in the 1970s, while neither can account for the data by itself. When we simply endow agents with cohort median inflation expectations from the Michigan Survey, our model predicts a portfolio shift towards housing very close to what we see in the data. However, it also generates counterfactually high credit volumes and nominal interest rate. Inflation uncertainty is a plausible feature of beliefs that makes trading in nominal credit less attractive. Our preferred specification for beliefs combines the two forces to account jointly for household wealth, portfolio weights, credit volumes and prices.

Our methodological approach differs from other studies of asset pricing with heterogeneous agents and may be interesting in its own right. Our goal is to characterize asset prices that are consistent with household optimization. We share this goal with conventional analysis of Euler equations pioneered by Hansen and Singleton (1982) that relates asset prices to the dynamics of consumption. The difference is that we relate asset prices to the distribution of household characteristics and the supply from other sectors, trading period by trading period. However, like Hansen and Singleton, we examine a set of relationships between observables that must hold in general equilibrium as long as households solve our life cycle problem. Our theory thus has content even if we do not take a stand on the behavior of other sectors, such as the government, businesses and the rest of the world, much like Euler equation tests have content even if the dynamics of consumption are taken from the data.

A novel feature of our approach is that the distribution of the key endogenous state variable that drives individual savings and portfolio choice – financial wealth relative to permanent income – comes directly from the data. We can thus provide a new perspective on the role of heterogeneity for asset prices. Existing studies have shown that heterogeneity does not matter if asset demand functions are approximately linear over the support of a model-implied wealth distribution. We examine asset demand in the support of the *observed* wealth distribution. We find that while approximate aggregation

obtains *within* all but the youngest age cohorts, the distribution of wealth *across* cohorts matters a lot, as emphasized by Krueger and Kubler (2004).

The idea that demographics should matter for asset prices in theory has been discussed for a long time, although there is no consensus on its quantitative impact. Several studies have shown that demographics cannot fully account for changes in stock prices if equity is the only long-lived asset in the model (Abel, 2003, Geanakoplos, Magill and Quinzii, 2004). The effect of demographics on house prices has been explored by Mankiw and Weil (1989) as well as Ortalo-Magne and Rady (2006). A recent literature on the Great Stagnation argues that low real interest rates reflect the aging of baby boomers (for example, Gagnon, Johannsen and Lopez-Salido, 2016, or Eggertsson, Mehrotra and Robbins, 2019). These studies model the savings margin and the relationship between growth and real interest rates. Our focus is on portfolio choice and asset price fluctuations in a world with multiple long-lived assets that incur capital gains and losses.¹ With both equity and real estate present in nonzero net supply, demographic shifts impact aggregate savings and hence the wealth-GDP ratio, but they cannot account for the larger movements in the individual components of wealth, especially equity.

The effects of inflation on (real) asset prices is also a familiar theme in the literature. We build on earlier work that has stressed the interaction of taxes and inflation, in particular, Feldstein (1980), Summers (1981) and Poterba (1991). We evaluate the role of nominal rigidities in the tax code in a life cycle setup when agents disagree about inflation. Our assumptions on disagreement by age are consistent with Malmendier and Nagel (2016) who attribute it to households' updating beliefs over their lifetime. We show that disagreement interacts importantly with taxes: the fact that old households who hold most equity expected lower inflation reduces the strength of tax-driven portfolio shifts. We further emphasize that disagreement and inflation uncertainty have opposite effects on nominal credit volumes. The role of inflation uncertainty in the 1970s has been emphasized by Cogley and Sargent (2001).

There is a large body of work on quantitative asset pricing models with heterogeneous agents and incomplete markets. Its main focus is the properties of equity and bond returns in stationary rational expectations equilibria with real shocks, following the classic contributions of Constantinides and Duffie (1996), Heaton and Lucas (1996), and Krusell and Smith (1998). Constantinides et al. (2002), and Storesletten, Telmer and Yaron (2007) consider OLG models with equity, riskless bonds and uninsurable income risk. Favilukis, Ludvigson and van Nieuwerburgh (2017) also allow for housing in a production economy. Chien, Cole and Lustig (2012, 2016) study models with different perceptions and constraints, and hence disagreement about asset payoffs. We share with this literature the theme that movements of the wealth distribution can matter for prices, as the identity of the average investor changes. Our approach of measuring that distribution directly clarifies the strength of this mechanism. We further differ from this literature in our focus on the 1970s and the unique events that shaped it.

A recent literature studies the boom-bust episodes of the 2000s (see Piazzesi and Schneider, 2016 for

¹Portfolio choice with housing has been considered by Flavin and Yamashita (2002), Campbell and Cocco (2003), and Fernandez-Villaverde and Krueger (2011) Cocco (2005) and Yao and Zhang (2005) study intertemporal choice problems with three assets that are similar to the problem solved by our households.

a survey). Emphasis on high frequency price changes and housing policy leads many of these studies to introduce details of housing and mortgage markets, such as transaction costs, long term contracts, housing tenure, and heterogeneity of the housing stock. Our focus is different: since we want to study low frequency changes in asset allocation, our model is simpler along those dimensions, but explicitly incorporates portfolio choice between equity, housing, and risky nominal credit. Another shared theme with the literature is the supply of assets to the household sector. Justiniano, Primiceri and Tambalotti (2019) have shown that an increase in credit supply played an important role for interest rates and house prices in the early 2000s. We show that the stability of housing and equity supply together with the decline in bond supply generated a large role for demand shifters.

The rest of the paper is organized as follows. Section II presents the model. Section III describes the quantitative implementation and documents properties of the model inputs, that is, the joint distribution of asset endowments and income as well as asset supply. Section IV derives predictions for optimal household behavior and compares them to the data. Section V provides intuition for price determination and the role of heterogeneity for aggregates. Section VI derives predictions under baseline expectations and shows that they help understand the evolution of the wealth-GDP ratio. Section VII considers the effect of inflation on portfolio composition in the 1970s.

II Model

The model describes the household sector's planning and asset trading in a single time period t .

A. Households

Households enter the period with assets and debt accumulated earlier. During the period, they earn labor income, pay taxes, consume and buy assets. Labor income is affected by idiosyncratic income shocks. Households can invest in three types of assets: long-lived equity and real estate as well as short-lived nominal bonds. Households face two types of aggregate risk. They face aggregate growth risk through stock dividends and the aggregate components of their housing dividends and labor income. They also face aggregate inflation risk when borrowing or lending because there is no risk-free asset in real terms. Households also face idiosyncratic risk which affects the return on individual houses and labor income streams. There are only three assets, so markets are incomplete.

Planning Horizon

Consumers alive at time t differ by endowment of assets and numeraire good as well as by age. Differences in age are represented by age-specific planning horizons T and age-specific survival probabilities for the next period. We now describe the problem of a typical consumer with a planning horizon $T > 0$.

Preferences

Consumers care about two goods, housing services and other (non-housing) consumption which serves as the numeraire. A consumption bundle of s_t units of housing services and c_t units of numeraire

yields utility

$$(1) \quad C_t = c_t^\delta s_t^{1-\delta}.$$

Preferences over (random) streams of consumption bundles $\{C_t\}$ are represented by the recursive utility specification of Epstein and Zin (1989). Utility at time t is defined as

$$(2) \quad U_t = \left(C_t^{1-1/\sigma} + \beta E_t \left[U_{t+1}^{1-\gamma} \right]^{\frac{1-1/\sigma}{1-\gamma}} \right)^{\frac{1}{1-1/\sigma}},$$

where $U_{t+T} = C_{t+T}$. Here σ is the intertemporal elasticity of substitution, γ is the coefficient of relative risk aversion towards timeless gambles and β is the discount factor. The expectation operator takes into account that the agent will reach the next period only with an age-specific survival probability.

Equity

Shares of equity can be thought of as trees that yield some quantity of numeraire good as dividend. A consumer enters period t with an endowment of $\bar{\theta}_t^e \geq 0$ units of trees. Trees trade in the equity market at the ex-dividend price p_t^e ; they cannot be sold short. A tree pays d_t^e units of dividend at date t . We summarize consumers' expectations about prices and dividends beyond period t by specifying expectations about returns. In particular, we assume that consumers expect to earn a (random) real return $R_{\tau+1}^e$ by holding equity between any two periods τ and $\tau + 1$, where $\tau \geq t$.

Real Estate

Real estate – or houses – may be thought of as trees that yield housing services. A consumer enters period t with an endowment of $\bar{\theta}_t^h \geq 0$ units of houses. Houses trade at the ex-dividend price p_t^h ; they cannot be sold short. To fix units, we assume that one unit of real estate (also referred to as one house) yields one unit of housing services at date t . There is a perfect rental market, where housing services can be rented at the price p_t^s . Moreover, every house requires a maintenance cost of m units of numeraire. If a consumer buys one unit of real estate, he obtains a dividend $d_t^h := p_t^s - m$. Consumers form expectations about future returns on housing and rental prices $\{R_\tau^h, p_\tau^s\}_{\tau > t}$.

Borrowing and Lending

Consumers can borrow or lend by buying or selling one period discount nominal bonds. A consumer enters period t with an endowment of \bar{b}_t units of numeraire that is due to past borrowing and lending in the credit market. In particular, \bar{b}_t is negative if the consumer has been a net borrower in the past. In period t , the consumer can buy or sell bonds at a price q_t . A consumer expects every bond bought to pay $1/\pi_{t+1}$ units of the numeraire in period $t + 1$. Here π_{t+1} is random and may be thought of as the expected change in the dollar price of numeraire. This is a simple way to capture that debt is typically denominated in dollars.² For every bond sold, the consumer expects to repay $(1 + \xi)/\pi_{t+1}$

²To see why, consider a nominal bond which costs q_t dollars today and pays of \$1 tomorrow, or $1/p_{t+1}^c$ units of numeraire consumption. Now consider a portfolio of p_t^c nominal bonds. The price of the portfolio is q_t units of numeraire and its payoff is $p_t^c/p_{t+1}^c = 1/\pi_{t+1}$ units of numeraire tomorrow. The model thus determines the price q_t of a nominal bond in dollars.

units of numeraire in period $t + 1$, where $\zeta > 0$ is an exogenous credit spread.³ Bond sellers – borrowers – face a collateral constraint: the value of bonds sold may not exceed a fraction ϕ of the value of all real estate owned by the consumer. For periods $\tau > t$, consumers form expectations about the (random) real return on bonds $\{R_\tau^b\}_{\tau>t}$, where $R_\tau^b = 1/q_{\tau-1}\pi_\tau$ is the (ex post) real lending rate, and $R_\tau^b(1 + \zeta)$ is the (ex post) real borrowing rate.

Non-Asset Income

Consumers are endowed with an age-dependent stream of numeraire good $\{y_\tau\}_{\tau=t}^{t+T}$. Here income should be interpreted as the sum of labor income, transfer income, and income on illiquid assets such as private businesses.

Budget Set

The consumer enters period t with an endowment of houses and equity $(\bar{\theta}_t^h, \bar{\theta}_t^e)$ as well as an endowment of $y_t + \bar{b}_t$ from non-asset income and past credit market activity. At period t prices, initial wealth is therefore

$$(3) \quad \bar{w}_t = (p_t^h + d_t^h)\bar{\theta}_t^h + (p_t^e + d_t^e)\bar{\theta}_t^e + \bar{b}_t + y_t.$$

To allocate this initial wealth to consumption and purchases of assets, the consumer chooses a plan $a_t = \{c_t, s_t, \theta_t^h, \theta_t^e, b_t^+, b_t^-\}$, where $b_t^+ \geq 0$ and $b_t^- \geq 0$ denote the amount of bonds bought and sold, respectively. It never makes sense for a consumer to borrow and lend simultaneously, that is, $b_t^+ > 0$ implies $b_t^- = 0$ and vice versa.

The plan a_t must satisfy the collateral constraint $q_t b_t^- \leq \phi p_t^h \theta_t^h$. The plan must also satisfy the budget constraint

$$(4) \quad c_t + p_t^s s_t + w_t = \bar{w}_t,$$

where terminal wealth is defined as

$$w_t = p_t^h \theta_t^h + p_t^e \theta_t^e + q_t b_t^+ - q_t b_t^-.$$

To formulate the budget constraint for periods beyond t , it is helpful to define the value of the consumer's stock portfolio in t by $w_t^e = p_t^e \theta_t^e$, the consumer's real estate portfolio by $w_t^h = p_t^h \theta_t^h$ as well as the values of a (positive or negative) bond portfolio, $w_t^{b+} = q_t b_t^+$ and $w_t^{b-} = q_t b_t^-$. For periods $\tau > t$, the consumer chooses plans $a_\tau = \{c_\tau, s_\tau, w_\tau^h, w_\tau^e, w_\tau^{b+}, w_\tau^{b-}\}$ subject to the collateral constraint

³One way to think about the organization of the credit market is that there is a financial intermediary that matches buyers and sellers in period t . In period $t + 1$, the intermediary will collect $(1 + \zeta) / \pi_{t+1}$ units of numeraire from every borrower (bond seller), but pay only $1 / \pi_{t+1}$ to every lender (bond buyer), keeping ζ / π_{t+1} for itself. We do not model the financial intermediary explicitly since we only clear markets in period t .

$w_\tau^{b-} \leq \phi w_\tau^h$ and the budget constraint

$$(5) \quad \begin{aligned} & c_\tau + p_\tau^s s_\tau + w_\tau^h + w_\tau^e + w_\tau^{b+} - w_\tau^{b-} \\ & = R_\tau^h w_{\tau-1}^h + R_\tau^e w_{\tau-1}^e + R_\tau^b w_{\tau-1}^{b+} - R_\tau^b (1 + \zeta) w_{\tau-1}^{b-} + y_\tau. \end{aligned}$$

We denote the consumer's overall plan by $a = (a_t, \{a_\tau\}_{\tau=t+1}^{t+T})$. This plan is selected to maximize utility (2) subject to the budget constraints (3)-(5) and the collateral constraints.

Taxes

In some of our examples below, we will assume proportional income taxes as well as capital gains and dividend taxes. This will not change the structure of the consumer's problem, just the interpretation of the symbols. In particular, labor income, dividends and returns will have to be interpreted as their after-tax counterparts. We will discuss their precise form in the calibration section below.

Terminal Consumers

The consumers described so far have planning horizons $T > 0$. We also allow consumers with planning horizon $T = 0$. These consumers also enter period t with asset and numeraire endowments that provide them with initial wealth \bar{w}_t , as in (3). However, they do not make any savings or portfolio decisions. Instead, they simply purchase numeraire and housing services in period t to maximize (1) subject to the budget constraint

$$c_t + p_t^s s_t = \bar{w}_t.$$

B. Equilibrium

To capture consumer heterogeneity, we assume a finite number of consumer types, indexed by i , with different initial endowment vectors $(\bar{\theta}_t^h(i), \bar{\theta}_t^e(i), y_t(i) + \bar{b}_t(i))$ and planning horizons $T(i)$.

New Asset Supply from the Rest of the Economy

To close the model and regulate the new supply of assets (which is exogenous to the household sector), we introduce a rest-of-the-economy (ROE) sector. It may be thought of as a consolidation of the business sector, the government and the rest of the world. The ROE sector is endowed with f_t^e trees and f_t^h houses in period t . A positive f_t^e represents the issuance of shares by the corporate sector, otherwise share repurchases. Similarly, a positive f_t^h captures new construction of houses. In addition, the ROE enters period t with an outstanding debt of \bar{B}_t units of numeraire, and it raises D_t units of numeraire by borrowing in period t . The surplus from these activities is

$$C_t^{ROE} = f_t^h (d_t^h + p_t^h) + f_t^e (d_t^e + p_t^e) + D_t - \bar{B}_t.$$

The ROE consumes the surplus C_t^{ROE} . More generally, the ROE has "deep pockets" out of which it pays for a deficit, $C_t^{ROE} < 0$.

Aggregate Asset Supply

We normalize initial endowments of equity and real estate such that there is a single tree and a single house outstanding:

$$\sum_i \bar{\theta}_t^h(i) = \sum_i \bar{\theta}_t^e(i) = 1.$$

In addition, we assume that initial endowments from past credit market activity are consistent, in the sense that every position corresponds to some offsetting position, either by a household or by the ROE sector:

$$\sum_i \bar{b}_t(i) = \bar{B}_t.$$

Equilibrium

An equilibrium consists of a vector of prices for period t , $(p_t^h, p_t^e, q_t, p_t^s)$, a surplus for the ROE sector C_t^{ROE} , as well as a collection of consumer plans for period t , $\{a_t(i)\} = \{c_t(i), s_t(i), \theta_t^h(i), \theta_t^e(i), b_t^+(i), b_t^-(i)\}$ such that

- (1) for every consumer, the plan $a_t(i)$ is part of an optimal plan $a(i) = (a_t(i), \{a_\tau(i)\}_{\tau=t+1}^{t+T(i)})$ given consumer i 's endowment, planning horizon, and expectations about future prices and returns;
- (2) markets for all assets and goods clear:

$$\begin{aligned}\sum_i \theta_t^h(i) &= 1 + f_t^h, \\ \sum_i \theta_t^e(i) &= 1 + f_t^e, \\ q_t \sum_i b_t^+(i) &= D_t + q_t \sum_i b_t^-(i), \\ \sum_i c_t(i) + m \theta_t^h(i) + C_t^{ROE} &= \sum_i y_t(i) + d_t^e (1 + f_t^e), \\ \sum_i s_t(i) &= \sum_i \theta_t^h(i).\end{aligned}$$

In addition to market clearing conditions for stocks, bonds and numeraire, there are two market clearing conditions for housing: one for the asset “real estate” and one for the good “housing services”. The first equation ensures that the total demand for houses equals their total supply. The fifth equation ensures that the fraction of houses that owners set aside as investment real estate – that is, selling services in the rental market – is the same as the fraction of housing services demanded in the rental market. As is common in competitive models, one of the five market clearing conditions is redundant, as it is implied by the sum of consumers’ budget constraints, the definition of C_t^{ROE} and the other four market clearing conditions. Solving for equilibrium prices thus amounts to solving a system of four equation in the four prices p_t^h, p_t^e, q_t and p_t^s .

C. Discussion of the Assumptions

Connection to National Income Accounting

Introducing the ROE sector allows us to accommodate deviations between consumption and house-

hold sector income that are usually ignored in endowment economy models. To clarify the connection to the FAUS/NIPA framework, we sum up the last three market clearing conditions and rearrange to obtain

$$\underbrace{\sum_i c_t(i) + p_t^s s_t(i)}_{\text{personal consumption}} + \underbrace{f_t^h p_t^h + f_t^e p_t^e + D_t - q_{t-1} \bar{B}_t}_{\text{personal savings}} = \underbrace{\sum_i y_t(i) + d_t^e + d_t^h + (1 - q_{t-1}) \bar{B}_t}_{\text{personal income}}$$

The first term on the left-hand side of this equation is personal consumption, including housing services. The right-hand side of this equation represents personal income. While our definition of income differs from the NIPAs in some details, as discussed in the next section, the basic components are the same: non-asset income, dividends on the two long-lived assets, and net interest. The difference between personal consumption and income is the second term on the left-hand side, which represents personal savings. It consists of the same components as in the FAUS: net acquisition of real estate and net acquisition of financial assets, here equity and bonds.⁴ The supply of assets provided by the ROE thus allows for positive or negative personal savings in equilibrium.

Asset Supply and Savings

The possibility of nonzero savings makes our model compatible with richer models that explicitly consider production, fiscal policy, and the fact that the US is not a closed economy. A richer model would give rise to explicit policy functions for the business, government, and foreign sectors that link those sectors' net asset supplies to market prices. At the same time, if the richer model accounts for observed supply by the ROE sector to the household sector, evaluating the policy functions at observed prices should deliver observed quantities. In our empirical implementation below, we thus set the net supply by the ROE equal to observed quantities.

We thus determine asset prices such that households are willing to hold observed quantities, which would also be provided by the ROE in a richer model that fits observed quantities. When the other side of the market is also derived from optimal policies, the household sector could be modelled in exactly the same way as in our model.⁵ Our approach thus avoids a common criticism of asset pricing models based on endowment economies: model features that help explain prices in an endowment economy – for example, certain preference specifications – may entail undesirable outcomes once savings are allowed. Our strategy matches simultaneously prices and observed household savings.

⁴FAUS savings is larger than our concept of savings because FAUS investment contains investment in noncorporate business, as well as purchases of consumer durables. In our empirical implementation, noncorporate business is treated as illiquid and investment in it is subtracted from income. In addition, we follow the NIPA convention of treating expenditure on durables as consumption. This seems appropriate given the 6-year length of each period in our model.

⁵Our modelling strategy requires some care in interpreting counterfactuals. For example, computing the response to a change in tax rates cannot tell us what would have happened to the economy as a whole, since we do not model supply responses to tax changes. Instead, the counterfactual tells us what the household sector would have done if taxes changed, and what the effect on prices would have been had supply remained unchanged. The reason the exercise is still useful is that it allows us to see whether tax changes can play a first order role in accounting for asset prices over time, given the actual movements in equilibrium quantities.

Prices vs. Quantities

It is not common in asset pricing models to distinguish asset prices and quantities. Indeed, if the dividend from a single “tree” provides all consumption, there is no meaningful concept of quantity – doubling the quantity of trees is the same as doubling the dividend on the single tree. In our model, asset quantity is identified because we allow trades between the household sector and the ROE. Households increase or decrease the quantity of trees they own *only* through such trades. The change in quantity is determined by comparing the change in market capitalization with the capital gain on pre-existing trees. While a capital gain is enjoyed by households who enter the periods owning trees, a change in capitalization requires households to pay for new trees.

Expectations

Our approach does not place restrictions on prices beyond date t itself. In particular, it does not require that agents agree on a common model structure that links future fundamentals and prices. As far as household behavior is concerned, our empirical implementation instead follows the literature on portfolio choice, where exogenous processes for returns and income are standard. However, we go beyond portfolio choice models in that we explore how equilibrium *prices* change when investors’ expectations about *returns* vary in a controlled way. This is particularly useful if expectations are matched to inflation surveys, as we do in Section VII, since expected inflation directly affects returns.

Our model does not link price expectations to future model-implied prices. Such a connection is imposed in rational expectations models, and more generally whenever agents have “structural knowledge” of the price function, such as in many models with Bayesian learning. Assuming structural knowledge has appeal when agents understand well how the economy works, for example because there are stable recurrent patterns in the data. Indeed, in a stationary environment, “knowing the price function” can be justified by learning from past data, and appears no more restrictive than “knowing the distribution of fundamentals”.

However, our interest in this paper is in low frequency asset-price movements during times of structural change. In this context, it is not clear why agent beliefs should be based on the particular model structure that we consider today with the benefit of hindsight. For example, it is not clear whether household expectations should be required to anticipate structural change that we only now know to have taken place. Since the usual justifications for narrowing down the set of expectations by imposing rational expectations or structural knowledge is not compelling, we turn to flexible modelling of exogenous expectations.

Connection to consumption-based asset-pricing

Many studies that examine asset prices jointly with real variables focus on consumption, in particular the connection between consumption and asset prices implied by intertemporal Euler equations. Since our model is based on utility maximization, an Euler equation also holds, at least at the individual household level. Instead of considering a map from endowments and supply to prices, one could consider examining the Euler equation directly. In certain ways this might be easier, since it requires

only one non-price variable, consumption, although serious treatment of heterogeneous agents would require household-level consumption data.

The main problem with the Euler equation approach in our context is that the Euler equation relates prices to investors' *planned* consumption, and thus makes sense only if planned consumption can be measured. Imposing rational expectations effectively makes the (subjective) distribution of planned consumption measurable by setting it equal – by assumption – to the observed distribution of realized consumption. As discussed above, this is a strong assumption that is not compelling in the context of structural change. As a result, the Euler equation approach itself is not appealing for the application we consider in this paper.

TABLE 1: NEGATIVE NET WORTH HOUSEHOLDS

	1962 SCF				1995 SCF				2001 SCF			
	avg.	29	53	77	avg.	29	53	77	avg.	29	53	77
% of households	4%	10%	3%	2%	7%	18%	4%	2%	7%	20%	5%	2%
net worth (in \$)	-.4K	-.4K	-.8K	-1.3K	-9K	-9K	-13K	-8K	-11K	-10K	-9K	-1K

NOTE: Row 1 reports the percentage of households with negative net worth in the U.S. population based on different SCFs. Row 2 reports the average net worth of these households in dollars.

Nonnegative Net Worth

Few households have negative net worth. Table 1 documents that the percentage of negative net worth households has always been between 4% and 7%. Table 1 also shows that the net worth of these households is moderate. For example, the average net worth was -11K Dollars in 2001. These numbers suggest that the most important reason for household borrowing is not consumption smoothing. Instead, young households “borrow to gamble” — they borrow to be able to buy more risky assets, such as housing.

The Role of Housing

In both our model and in reality, housing plays a dual role: housing can be used for consumption and investment. In its role as a consumption good, housing is different from other consumption, because it enters the utility function separately and has a different price p_t^s . In its role as asset, housing is different from other assets, because of its return properties (which will be discussed later), and because it can serve as collateral. In reality, the two roles are often connected, because the amount of housing services consumed is equal to the dividend paid by the amount of housing held in the portfolio. In our model, we abstract from this issue (at least for now) for several reasons. First, most homeowners not only own their primary residence, but also some investment real estate (such as time shares, vacation homes, secondary homes, etc.) Hence, there is no tight link between consumption and investment for these households. Second, U.S. households are highly mobile. As a consequence, houses

are owned by the same owner on average for only seven years, which is roughly the same length as a period in our model. Moreover, the decision to move is often more or less exogenous to households (for example, because of job loss, divorce, death of the spouse, etc.)

III Quantitative Implementation

The inputs needed for implementing the model are (i) the joint distribution of asset endowments and income, (ii) aggregate supply of assets to the household sector from other sectors, (iii) expectations about labor income and asset returns and (iv) parameter values for preferences and the credit market. We describe our choices in detail in the Appendix; here we provide a brief overview. We start with a description of what a model period corresponds to in the data.

Timing

The length of a period is six years. Since the model compresses what happens over a six year span into a single date, prices and holdings are best thought of as period averages. We assume that consumers expect to live for at most 10 such periods, where the first period of life corresponds roughly to the beginning of working life. In any given period, we consider 11 age groups of households (<23, 24-29, 30-35, 36-41, 42-47, 48-53, 54-59, 60-65, 66-71, 72-77, >77) who make portfolio choice decisions. For ease of comparison with other models, we nevertheless report numbers at annual rates.

In the time series dimension, we focus on three six-year periods, namely 1965-70, 1975-80, and 1992-97. To construct asset endowment distributions, we use data on asset holdings from the respective precursor periods 1959-64, 1969-74, and 1986-91. Micro data is not available at high frequencies. To capture the wealth and income distribution during a period, we have chosen the above intervals so that every period that contains a Survey of Consumer Finances contains is in the 4th year of the period; we thus use surveys from 1962, 1989 and 1995. We use the 4th year survey to infer income and asset holdings where possible.

A. Defining Assets and Income

We map the three assets in the model to three broad asset classes in US aggregate and household statistics. For equity and real estate, we use the Financial Accounts of the United States (FAUS) to derive measures of (i) aggregate holdings of the household sector, (ii) net purchases by the household sector and (iii) aggregate dividends received by the household sector. We use the numbers on value, dividends and new issues to calculate price dividend ratios and holding returns on equity and real estate. For both assets, we also define corresponding measures at the individual level using the Survey of Consumer Finances (SCF).

Equity

We identify *equity* with shares in corporations held and controlled by households. We include both publicly traded and closely held shares, and both foreign and domestic equity. We also include shares held indirectly through investment intermediaries if the household can be assumed to control the asset

allocation (into our broad asset classes) himself. We take this to be true for mutual funds and defined contribution (DC) pension plans. We do not include equity in defined benefit (DB) pension plans, since households typically do not control the asset allocation of such funds. Our concept of dividends thus equals dividends received by households from the National Income and Product Accounts (NIPA) less dividends on household holdings in DB plans.

Real estate

Our concept of *residential real estate* contains owner-occupied housing, directly held residential investment real estate, as well as residential real estate recorded in the FAUS/NIPA as held indirectly by households through noncorporate businesses. This concept contains almost all residential real estate holdings, since very few residential properties are owned by corporations. We take housing dividends to be housing consumption net of maintenance and property tax from NIPA. For net purchases of new houses, we use aggregate residential investment from NIPA.

The ROE sector endowment of equity consists of net new equity purchased by the household sector during the trading period. The factor f^e states this endowment relative to total market capitalization in the model. We thus use net purchases of corporate equity divided by total household holdings of corporate equity. The ROE sector endowment of residential real estate consists of residential investment, divided by the value of residential real estate. The top panel of Figure 3 plots quarterly series for both ROE endowment, or “asset supply” numbers. The calibration of the model uses six-year averages of these series.

Nominal positions

Our concept for a household’s *bond* holdings is its net nominal position, that is, the market value of all nominal assets minus the market value of nominal liabilities. As for equity, holdings include not only direct holdings, but also indirect holdings through investment intermediaries. To calculate market value, we use the market value adjustment factors for nominal positions in the U.S. from Doepke and Schneider (2006). In line with our treatment of tenant-occupied real estate, we assign residential mortgages issued by noncorporate businesses directly to households.

When considering *gross* nominal positions, we must take into account the fact that some netting of positions occurs at the individual level. A typical household will have both a mortgage loan or credit card and a savings account or bonds in a pension fund. However, the model has only one type of nominal asset, and a household can be either long or short that asset. If we were to match the gross aggregates from the FAUS or SCF in our model, this would inevitably lead to net positions that are too large. Instead, we sort SCF households into borrowers and lenders, according to whether their net nominal position is negative or not. The numbers for gross borrowing and lending are then calculated as minus the sum of net nominal positions of borrowers as well as the sum of net positions of all lenders, respectively.

Table 2 summarizes these gross nominal positions after individual netting from the SCF and compares them to those in the FAUS. Both the FAUS numbers and our estimates reflect a steady increase in

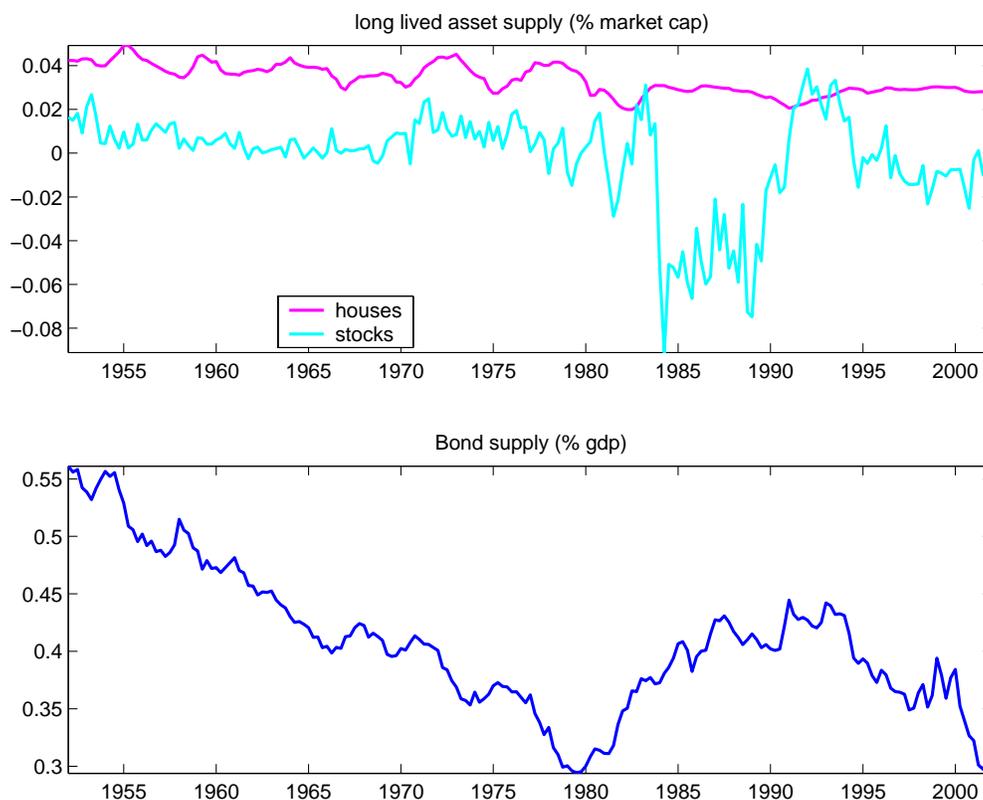


Figure 3: *Top panel:* Net new corporate equity as a percent of total household holdings of corporate equity and residential investment as percent of value of residential real estate. Data are quarterly at annual rates. *Bottom panel:* Net nominal position of household sector as a percent of GDP.

borrowing by the household sector. At the same time, both sets of numbers show a reduction in nominal asset holdings in the 1970s followed by an increase between 1978 and 1995. Throughout, individual netting reduces gross lending by roughly one third, while it reduces gross borrowing by slightly more than half.

TABLE 2: GROSS BORROWING AND LENDING (%GDP)

	1968		1978		1995	
	lending	borrowing	lending	borrowing	lending	borrowing
FAUS aggregates	88	47	84	51	107	68
SCF after indiv. netting	61	20	56	23	70	31

The initial nominal position of the ROE sector is taken to be minus the aggregate (updated) net nominal position of the household sector. Finally, the new net nominal position of the ROE in period t – in other words, the “supply of bonds” to the household sector – is taken to be minus the aggregate net nominal positions from the FAUS for period t . This series is reproduced in the bottom panel of Figure 3.

Our concept of non-asset *income* comprises all income that is available for consumption or investment, but not received from payoffs of one of our three assets. We construct an aggregate measure of such income from NIPA and then derive a counterpart at the household level from the SCF. Of the various components of worker compensation, we include only wages and salaries, as well as employer contributions to DC pension plans. We do not include employer contributions to DB pension plans or health insurance, since these funds are not available for consumption or investment. However, we do include benefits disbursed from DB plans and health plans. Also included are transfers from the government. Finally, we subtract personal income tax on non-asset income.

B. The joint distribution of asset endowments and income

Consumers in our model are endowed with both assets and non-asset income. To capture decisions made by the cross-section of households, we thus have to initialize the model for every period t with a *joint* distribution of asset endowment and income. We derive this distribution from data on terminal asset holdings and income in the precursor period $t - 1$. To handle multidimensional distributions, we approximate them by a finite number of household types. Types are selected to retain key moments of the full distribution, in particular aggregate gross borrowing and lending.

Since the aggregate endowment of long-lived assets is normalized to one, we can read off the endowment of a household type in period t from its *market share* in period $t - 1$. For each long-lived asset $a = h, e$, suppose that $w_{t-1}^a(i)$ is the market value of investor i 's position in $t - 1$ in asset a . Its initial holdings are given by

$$\begin{aligned}\bar{\theta}_t^a(i) &= \theta_{t-1}^a(i) = \frac{w_{t-1}^a(i)}{\sum_i w_{t-1}^a(i)} = \frac{p_{t-1}^a \theta_{t-1}^a(i)}{p_{t-1}^a \sum_i \theta_{t-1}^a(i)} \\ &= \text{market share of household } i \text{ in period } t - 1.\end{aligned}$$

For nominal assets, the above approach does not work since these assets are short-term in our model. Instead, we determine the market value of nominal positions in period $t - 1$ and update it to period t by multiplying it with a nominal interest rate factor:

$$\begin{aligned}\bar{b}_t(i) &= (1 + i_{t-1}) \frac{w_{t-1}^b(i)}{\text{GDP}_t} \\ &= (1 + i_{t-1}) \frac{w_{t-1}^b(i)}{\sum_i w_{t-1}^b(i)} \frac{\sum_i w_{t-1}^b(i)}{\text{GDP}_{t-1}} \frac{\text{GDP}_{t-1}}{\text{GDP}_t}.\end{aligned}$$

Letting g_t denote real GDP growth and D_t the aggregate net nominal position as a fraction of GDP, we have

$$\bar{b}_t(i) \approx \bar{\theta}_{t-1}^b(i) D_{t-1} (1 + i_{t-1} - g_t - \pi_t).$$

This equation distinguishes three reasons why $\bar{b}_t(i)$ might be small in a given period. The first is simply

that the household's nominal investment in the previous period was small. Since all endowments are stated relative to GDP, all current initial nominal positions are also small if the economy has just undergone a period of rapid growth. Finally, initial nominal positions are affected by surprise inflation over the last few years. If the nominal interest rate i_{t-1} does not compensate for realized inflation π_t , then \bar{b}_t^i is small (in absolute value). Surprise inflation thus increases the negative position of a borrower, while it decreases the positive position of a borrower.

The final step in our construction of the joint income and endowment distribution is to specify the marginal distribution of non-asset income. Here we make use of the fact that income is observed in period $t - 1$ in the SCF. We then assume that the transition between $t - 1$ and t is determined by a stochastic process for non-asset income. We employ the same process that agents in the model use to forecast their non-asset income, described in Appendix B. This approach allows us to capture the correlation between income and *initial* asset holdings that is implied by the joint distribution of income and wealth.

Distributions for 1968, 1978 and 1995

Figure 4 provides summary information on asset endowment and income distributions in the three trading periods we consider below. The trading periods are identified in the figure by their respective fourth year: 1968, 1978 and 1995. The top left panel provides population weights by cohort. Cohorts are identified on the horizontal axis by the upper bound of the age range. In addition, the fraction of households that exit during the period are offset to the far right.

The different years can be distinguished by the line type: solid with circles for 1968, dashed with squares for 1978 and dotted with diamonds for 1995. Using the same symbols, the top right panel shows house endowments (light lines) and stock endowments (dark lines) by age cohort, while the bottom left panel shows initial net nominal positions as a percent of GDP. Finally, the bottom right panel shows income distributions. Here we plot not only non-asset income, but also initial wealth not invested in long-lived assets, in other words,

$$(6) \quad E_t = d_t^h \bar{\theta}_t^h + d_t^e \bar{\theta}_t^e + \bar{b}_t + y_t.$$

This aggregate will be useful to interpret the results below.

Two demographic changes are apparent from the figure. First, the baby boom makes the two youngest cohorts relatively larger in the 1978 cross section than in the other two years. By 1995, the boomers have aged so that the 42-47 year olds are now the largest cohort. This shift of population shares is also reflected in the distribution of income in the bottom right panel. Second, the relative size of the oldest group has become larger over time. Recently, a lot of retirement income comes from assets, so that the share of E_t of the elderly groups has also increased a lot. A key difference between the 1968 and 1978 distributions is thus that the latter places more weight on households who tend to save little: the oldest and, especially, the youngest. While the 1995 distribution also has relatively more weight on the elderly, it emphasizes more the middle-aged rather than the young.

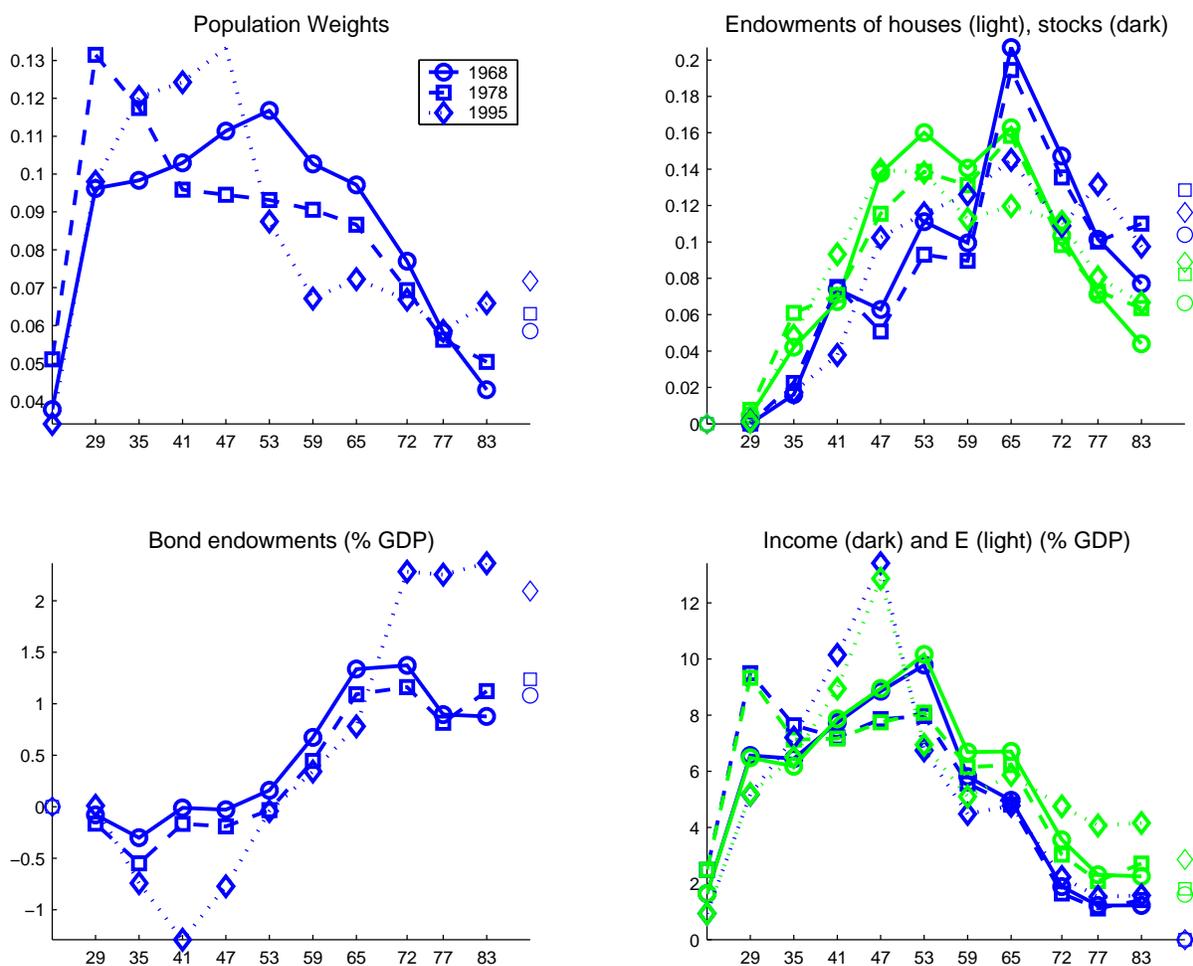


Figure 4: Asset endowment and income distributions in 1968, 1978 and 1995. *Top left panel:* Population weights by cohort, identified on the horizontal axis by the upper bound of the age range. Exiting households during the period are on the far right. *Top right panel:* House endowments (light lines) and stock endowments (dark lines) by age cohort. *Bottom left panel:* initial net nominal positions as a percent of GDP. *Bottom right panel:* Income distributions.

The comparison of stock and house endowments in the top right panel reveals that housing is more of an asset for younger people. For all years, the market shares of cohorts in their thirties and forties are larger for houses than for stocks, while the opposite is true for older cohorts. By and large, the market shares are however quite similar across years. In contrast, the behavior of net nominal positions relative to GDP (bottom right hand panel) has changed markedly over time. In particular, the amount of intergenerational borrowing and lending has increased: young households today borrow relatively more, while old households hold relatively more bonds.

C. Expectations and Preference Parameters

A baseline set of beliefs for returns and non-asset income is derived in Appendix B. We assume that consumers believe real asset returns and aggregate growth to be serially independent over successive six year periods. Moreover, consumers believe that returns and growth are identically distributed for periods beyond $t + 1$. However, in most of the exercises below, we allow beliefs for returns between t and $t + 1$ to differ from baseline beliefs, so that returns are not iid. (We discuss the latter aspect of beliefs below when we present our results.) To pick numbers, we start from empirical moments of six-year ex-post pre-tax real returns on fixed income securities, residential real estate and equity, as well as inflation and growth. Since returns on individual properties are more volatile than those on a nationwide housing index, we add an idiosyncratic shock to the housing return faced by an individual household.⁶

We also specify a stochastic process for after-tax income. Briefly, the functional form for this process is motivated by existing specifications for labor income that employ a deterministic trend to capture age-specific changes in income, as well as permanent and transitory components. We also use estimates from the literature to account for changes in the volatility of the different components over time.

As for preference parameters, the intertemporal elasticity of substitution is $\sigma = .5$, the coefficient of relative risk aversion is $\gamma = 5$ and the discount factor is $\beta = \exp(-0.025 \times 6)$. Since γ is low, agents do not want to hold bonds when faced simply with the point estimate of the historical Sharpe ratio on equity. We assume that agents take into account parameter uncertainty about the mean return on stocks and hence the equity premium. To capture this idea in a simple way, we reduce the mean return by a fixed spread.⁷ We select a mean adjustment on equity of 4.5% per year. With these numbers, we then match the *aggregate* portfolio weights reported in Table 3 almost exactly.

TABLE 3: 1995 DATA AND BASELINE BELIEFS

	Wealth/ GDP	Portfolio Share			Lend/ GDP	Borr/ GDP	PD Ratios		Nominal Rate
		Bond	Housing	Equity			Housing	Equity	
(1) —1995 data —	2.51	0.15	0.59	0.26	0.70	0.31	20.4	23.9	0.061
(2) baseline	2.52	0.15	0.60	0.25	0.70	0.31	20.9	23.2	0.062

NOTE: The first row reports the aggregate portfolio weights on bonds, housing and stocks from Panel (a) in Figure 2, the gross borrowing and lending numbers from Section A., the wealth-to-GDP ratio from Figure 1, and the price-dividend ratios for housing and stocks from Panel (b) in Figure 2 together with the nominal 6-year interest rate. The second row report the results computed from the model for 1995 with baseline beliefs.

⁶We set the idiosyncratic volatility so that the overall volatility of the return on an individual property is 11%, consistent with the estimates in Landvoigt, Piazzesi and Schneider (2015).

⁷Microfoundations for this approach are provided by Ilut and Schneider (2014) who consider a model with Knightian uncertainty. They capture uncertainty about the distribution of a particular shock by an interval of means; investors then evaluate portfolios by using the worst case mean. In particular, investors who go long in stocks behave *as if* they are pessimistic about stocks.

An alternative strategy would be to work with agents who have “high aversion against low perceived risk.” In this case, agents base their portfolio choice on the historical variances from Table B.2, but are characterized by high risk aversion, $\gamma = 25$, and high discounting, $\beta = \exp(-0.07 \times 6)$. The high γ is needed to lower the portfolio weight on bonds, while the low β is needed to reduce the precautionary savings motive in the presence of uninsurable income shocks. While the tables below report results based on agents with “low aversion against high perceived risk,” we would get results comparable to those in Table 3 based on this alternative parametrization.⁸

D. Taxes and the Credit Market

It remains to select parameters to capture taxes on investment as well as consumers’ opportunities for borrowing. We set the maximal loan-to-value ratio to $\phi = .8$, the modal value in the United States. We further choose spreads between borrowing and lending rates to match aggregate gross credit volumes in the economy. For 1995, we back out a value of 1.6%, in the ball park of existing estimates. Before the deregulation of banking and the increase in securitization of the early 1980s, mortgage markets were less developed. We back out the spread for the earlier episodes from credit volumes in 1968, and find a value of 2%.

Investors care about after-tax real returns. In particular, taxes affect the relative attractiveness of equity and real estate. On the one hand, dividends on owner-occupied housing are directly consumed and hence not taxed, while dividends on stocks are subject to income tax. On the other hand, capital gains on housing are more easily sheltered from taxes than capital gains on stocks. This is because many consumers simply live in their house for a long period of time and never realize the capital gains. Capital gains tax matters especially in inflationary times, because the nominal gain is taxed: the effective real after tax return on an asset subject to capital gains tax is therefore relatively lower with inflation. Inflation also increases the subsidy from mortgage deductability.

To measure the effect of capital gains taxes, one would ideally like to explicitly distinguish realized and unrealized capital gains. However, this would involve introducing state variables to keep track of past individual asset purchase decisions. To keep the problem manageable, we adopt a simpler approach: we adjust our benchmark returns to capture the effects described above. For our baseline set of results, we assume proportional taxes, and we set both the capital gains tax rate and the income tax rate to 20%. This number is consistent with the numbers reported by Sialm (2009) for average effective tax rates on equity returns as well as capital gains tax rates in the 1970s. We define after tax real stock returns by subtracting 20% from realized net real stock returns and then further subtracting 20% times the realized rate of inflation to capture the fact that nominal capital gains are taxed. In contrast, we assume that returns on real estate are not taxed.

⁸Yet another way to obtain realistic aggregate portfolio weights is to combine low risk aversion with first-time participation costs, as shown by Gomes and Michaelides (2005).

IV Household Behavior

In this section, we consider savings and portfolio choice in the cross section of households. We focus on baseline beliefs for 1995, the year we have used to calibrate beliefs. The initial distribution of asset endowments and income is derived from the 1989 SCF, as discussed in Section III. We first present optimal policies as functions of wealth and income. We then compare the predictions of the model for the cross-section of households in 1995 to actual observations from the 1995 SCF.

A. Lifecycle Savings and Portfolios

Since preferences are homothetic and all constraints are linear, the optimal savings rate and portfolio weights depend only on age and the ratio of initial wealth \bar{w}_t – defined in (3) above as asset wealth plus non-asset income – to the permanent component of non-asset income, say \hat{y}_t . For simplicity, we refer to \bar{w}_t/\hat{y}_t as the wealth-to-income ratio. Figure 5 plots agent decisions as a function of this wealth-to-income ratio.

Savings

The bottom right panel shows the ratio w_t/\bar{w}_t of terminal wealth to initial wealth, that is, the savings rate out of initial wealth. Savings are always positive, since the borrowing constraint precludes strategies that involve negative net worth. Investors who have more income in later periods than in the current period thus cannot shift that income forward by borrowing. In this sense, there is no borrowing for “consumption smoothing” purposes: all current consumption must instead come out of current income or from selling initial asset wealth.

If initial wealth is very low relative to income, all assets will be sold and all income consumed, so that the investor enters the next period with zero asset wealth. As initial wealth increases, a greater fraction of it is saved for future consumption. In the absence of labor income, our assumption of serially independent returns implies a constant optimal savings rate. As wealth becomes large relative to the permanent component of income, the savings rate converges to this constant.

The bottom right panel also illustrates how the savings rate changes with age. There are two relevant effects. On the one hand, younger investors have a longer planning horizon and therefore tend to spread any wealth they have over more remaining periods. This effect by itself tends to make younger investors save more. On the other hand, the non-asset income profile is hump-shaped, so that middle-aged investors can rely more on labor income for consumption than either young or old investors. This tends to make middle-aged investors save relatively more than other investors.

The first effect dominates when labor income is not very important, that is, when the wealth-to-income ratio is high. The figure shows that at high wealth-to-income ratios, the savings rate of the 29-35 year old group climbs beyond that of the oldest investor group. It eventually also climbs above the savings rate of the 48-53 year old group. The second effect is important for lower wealth-to-income ratios, especially in the empirically relevant range around 1-2, where most ratios lie in the data. In this region, the savings rate of the middle-aged is highest, whereas both the young and the old save less.

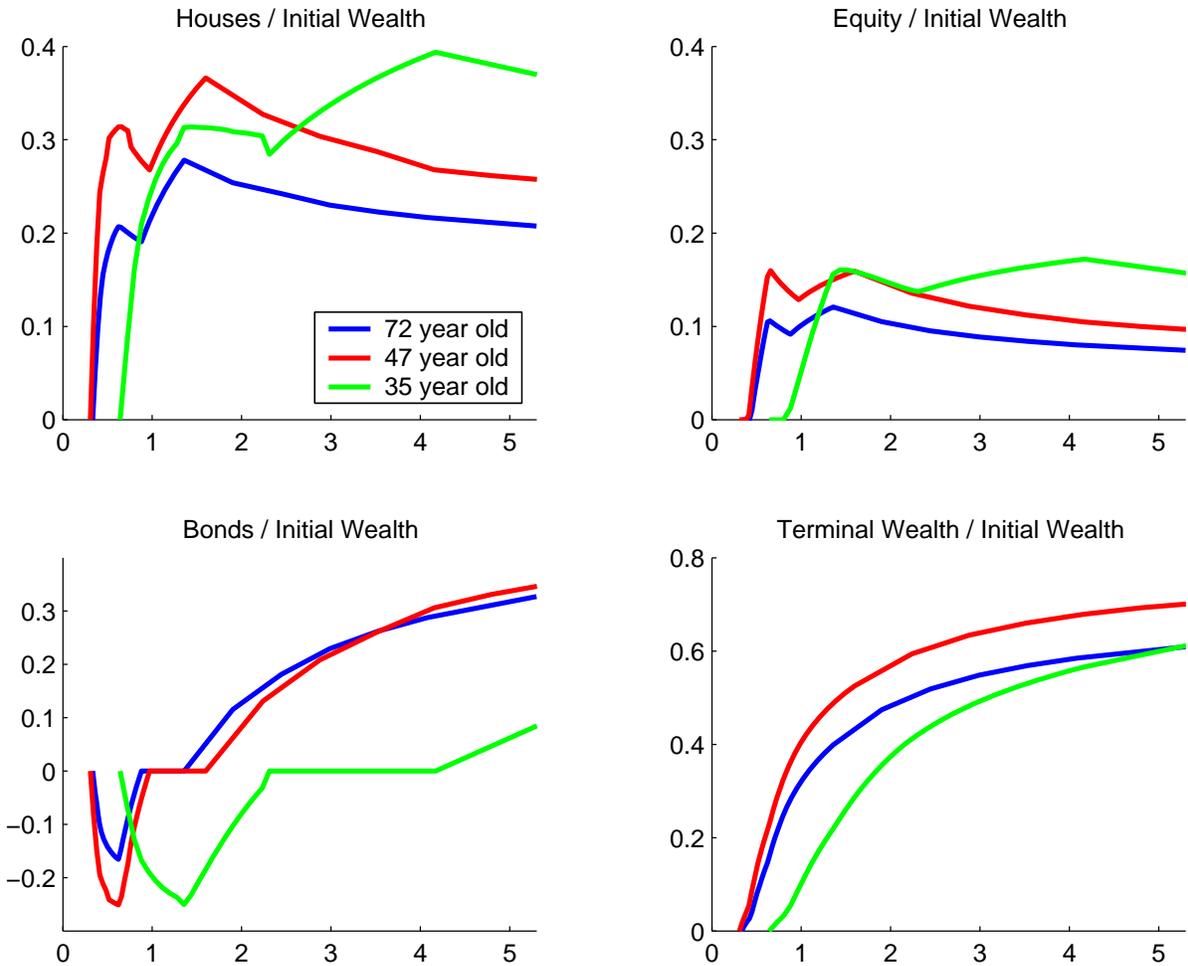


Figure 5: Asset holdings and terminal wealth, both as fractions of initial wealth, plotted against the initial wealth-to-income ratio. Age groups are identified by maximum age in the cohort.

Among the latter two groups, the young save the least when their wealth-to-income ratio is low.

Borrowing and Leverage

Rather than enable consumption smoothing, the role of borrowing in our model is to help households construct leveraged portfolios. The bottom left panel of Figure 5 shows that investors who are younger and have lower wealth-to-income ratios tend to go short in bonds. The top panels show that the borrowed funds are used to build leveraged positions of houses and also stocks. In contrast, investors who are older and have higher wealth-to-income ratios tend to go long in all three assets. Along the wealth-to-income axis, there is also an intermediate region where investors hold zero bonds. This region is due to the credit spread: there exist ratios where it is too costly to leverage at the high borrowing rate, while it is not profitable to invest at the lower lending rate.

The reason why “gambling” with leverage decreases with age and the wealth-to-income ratio is the presence of labor income. Effectively, an investor’s portfolio consists of both asset wealth and

human wealth. Younger and lower wealth-to-income households have relatively more human wealth. Moreover, the correlation of human wealth and asset wealth is small. As a result, households with a lot of labor income hold riskier strategies in the asset part of their portfolios. This effect has also been observed by Jagannathan and Kocherlakota (1996), Heaton and Lucas (2000), and Cocco (2005).

Stock v. House Ownership

For most age groups and wealth-to-income ratios, investment in houses is larger than investment in stocks. This reflects the higher Sharpe ratio of houses as well as the fact that houses serve as collateral while stocks do not. The latter feature also explains why the ratio of house to stock ownership is decreasing with both age and wealth-to-income ratio: for richer and older households, leverage is less important, and so the collateral value of a house is smaller.

Currently, the model cannot capture the fact that the portfolio weight on stocks tends to increase with the wealth-to-income ratio. While it is true in the model that people with higher wealth-to-income own more stocks relative to housing, they also hold much more bonds relative to both of the other assets. As a result, their overall portfolio weight on stocks actually falls with the wealth-to-income ratio. The behavior of the portfolio weight on stocks implies that the model produces typically too little concentration of stock ownership.

B. The Cross Section of Asset Holdings

Figure 6 plots predicted portfolio weights and market shares for various groups of households for 1995, given baseline beliefs. The panels also contain actual weights and market shares for the respective groups from the 1995 SCF. It is useful to compare both portfolio weights and market shares, since the latter also require the model to do a good job on savings behavior. Indeed, defining aggregate initial wealth $\bar{W} = \sum_i \bar{w}(i)$, the market share of asset $a = h, e$ for a household i can be written as

$$\theta^a(i) = \frac{\alpha^a(i) \bar{w}(i)}{\sum_i \alpha^a(i) \bar{w}(i)} = \frac{\alpha^a(i)}{\sum_i \alpha^a(i) \frac{\bar{w}(i)}{\bar{W}}} \frac{\bar{w}(i)}{\bar{W}} = \frac{\alpha^a(i)}{\bar{\alpha}^a} \frac{\bar{w}(i)}{\bar{W}},$$

where $\alpha^a(i)$ is household i 's portfolio weight and $\bar{\alpha}^a$ is the aggregate portfolio weight on asset a . A model that correctly predicts the cross section of portfolio shares will therefore only correctly predict the cross section of market shares if it also captures the cross section of terminal wealth. The latter in turn depends on the savings rate of different groups of agents.

The first row of Figure 6 documents savings behavior by cohort and wealth level. The top left panel plots terminal wealth as a fraction of GDP at the cohort level (blue/black lines) for the model (dotted line) and the data (solid line). It also shows separately terminal wealth of the top decile by net worth (green/light gray lines), again for the model and the data. This color coding of plots is maintained throughout the figure, so that a "good fit" means that the lines of the same color are close to each other.

The top left panel shows that the model does a fairly good job at matching terminal wealth. The model also captures skewness of the distribution of terminal wealth and how this skewness changes

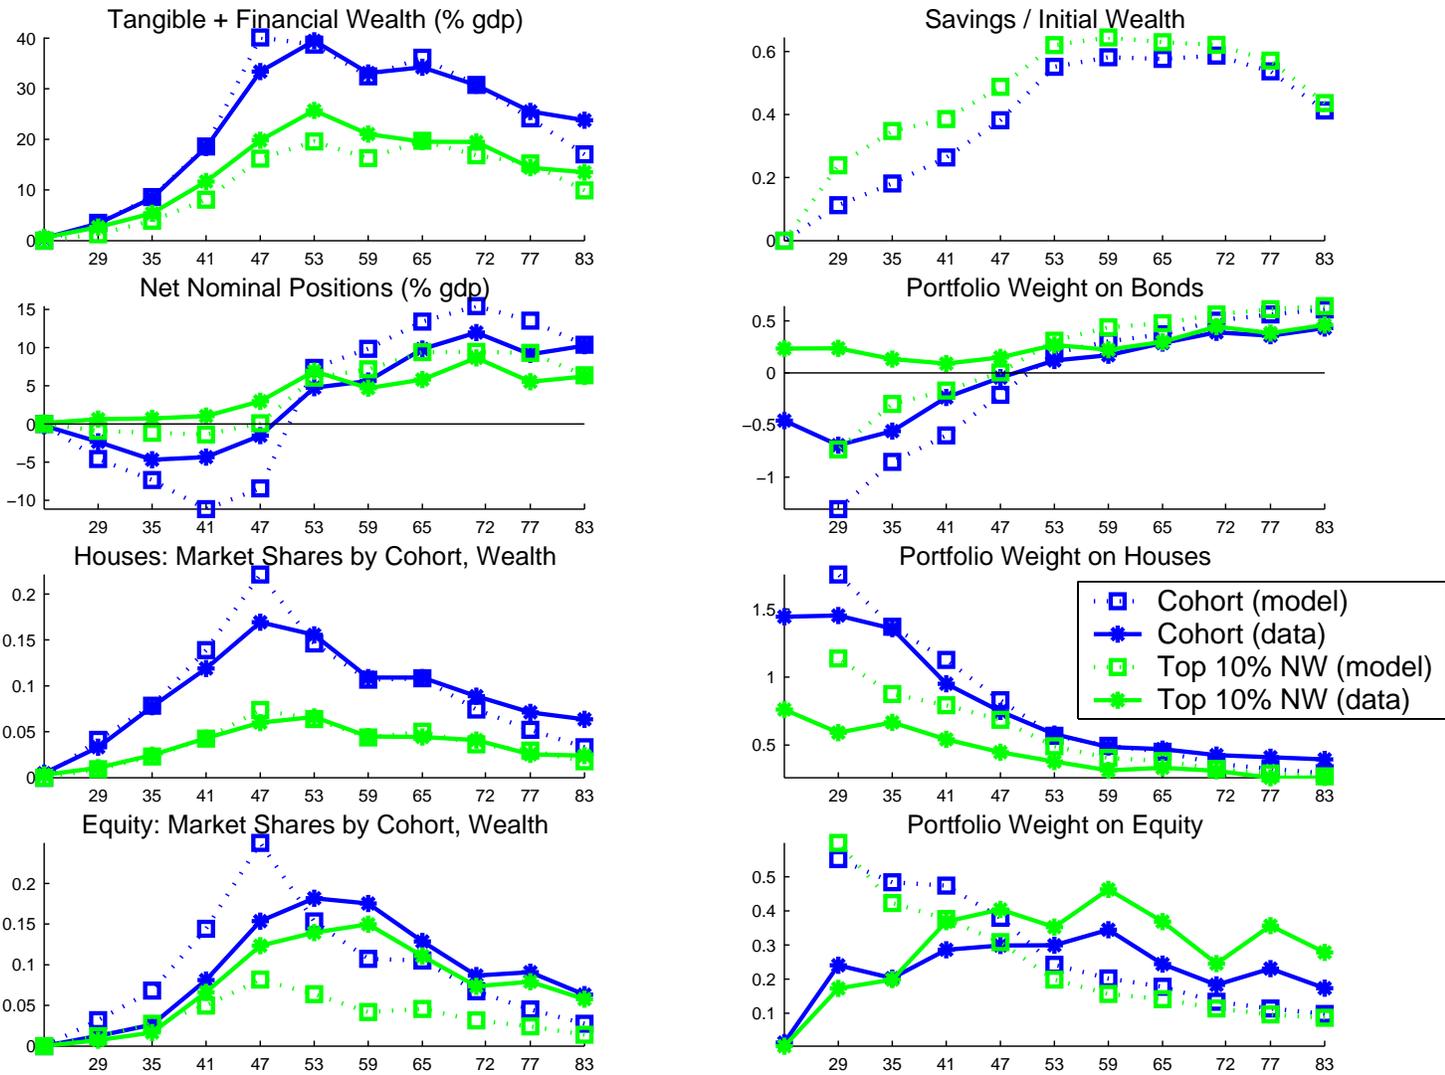


Figure 6: Market shares and portfolio weights for 1995

with age. The top 10% by net worth own more than half of total terminal wealth and their share is increasing with age. In the model, these properties are inherited in part from the distributions of endowment and labor income. However, it is also the case that richer agents save more out of initial wealth. This feature is apparent from the top right panel of Figure 6 which reports savings rates by cohort and net worth. It obtains because (i) the rich have higher ratios of initial wealth relative to current labor income, and (ii) the savings rate is increasing with the wealth-to-income ratio, as explained in the previous subsection.

In the data, the main difference in portfolio weights by age is the shift from houses into bonds over the course of the life cycle. This is documented in the right column of Figure 6. Young agents borrow in order to build leveraged positions in houses. In the second panel, their portfolio weights become

positive with age as they switch to being net lenders. The accumulation of bond portfolios makes houses – shown in the third panel on the right – relatively less important for older households. The model captures this portfolio shift fairly well. Intuitively, younger households “gamble” more, because the presence of future labor income makes them act in a more risk tolerant fashion in asset markets.

The left column shows the corresponding cohort aggregates. Nominal positions relative to GDP (second panel on the left) are first negative and decreasing with age, but subsequently turn around and increase with age so that they eventually become positive. These properties are present both in the model and the data. On the negative side, the model somewhat overstates heterogeneity in positions by age: there is too much borrowing – and too much investment in housing – by young agents. In particular, the portfolio weights for the very youngest cohort are too extreme. However, since the wealth of this cohort is not very large, its impact on aggregates and market shares is small.

For houses (third panel on the left), the combination of portfolio and savings choices generates a hump shape in market share. While younger agents have much higher portfolio weights on real estate than the middle-aged, their overall initial wealth is sufficiently low, so that their market share is lower than that of the middle-aged. Another feature of the data is that the portfolio shift from housing to bonds with age is much less pronounced for the rich. The model also captures this feature, as shown by the green/gray lines in the second and third rows of the figure. The intuition again comes from the link between leverage and the wealth-to-income ratio: the rich are relatively asset-rich (high wealth-to-income) and thus put together less risky asset portfolios, which implies less leverage and lower weights on housing.

The panels in the last row of Figure 6 plot market shares and portfolio weights for equity. This is where the model has the most problems replicating the SCF observations. Roughly, investment in stocks in the model behaves “too much” like investment in housing. Indeed, the portfolio weight is not only decreasing with age after age 53, as in the data, but it is also decreasing with age for younger households. As a result, while the model does produce a hump-shaped market share, the hump is too pronounced and occurs at too young an age. In addition, the model cannot capture the concentration of equity ownership in the data: the rich hold relatively too few stocks.

V Asset prices, Expectations, Endowments and Supply

In this section, we consider the map from model inputs – the distribution of endowments and income, asset supply, and expectations – into equilibrium prices. We report counterfactual experiments, based on the 1995 distribution, to illustrate the mechanics of the model and the relative magnitude of key effects.

Price Determination

To organize the discussion, we use the following equations, derived from the market clearing con-

ditions and the household budget constraints. Dropping time subscripts to simplify notation, we have

$$(7) \quad \begin{aligned} p^a (1 + f^a) &= \sum_i \alpha^a (\bar{w}(i) / \hat{y}(i); i, q) \bar{w}(i) ; \quad a = h, e. \\ p^h (1 + f^h) + p^e (1 + f^e) + D &= \sum_i S (\bar{w}(i) / \hat{y}(i); i, q) \bar{w}(i). \end{aligned}$$

The aggregate value of long-lived asset a (where $a = h$ or e) must equal, in equilibrium, the value of household demand for the asset. At the level of the individual household, asset demand can be written as the product of initial wealth $\bar{w}(i)$ and the asset weight α^a . The latter was shown in Figure 5 as a function of the wealth-to-income ratio $\bar{w}(i) / \hat{y}(i)$. It also depends on other household characteristics such as age and expectations, as well as on the market interest rate, or equivalently the bond price q .⁹ The second equation says that the value of all assets held by households – in our empirical framework, the numerator of the wealth-GDP ratio – must equal total household savings, where the savings rate, shown in the bottom right panel of Figure 5, is denoted S .

Equations (7) implicitly determine the three asset prices q , p^e and p^h as a function of supply, endowments and expectations, where the latter enter through the portfolio weight functions α^a and the savings rate function S . The second equation can be written as

$$\begin{aligned} p^h (1 + f^h) + p^e (1 + f^e) + D &= \sum_i \frac{\bar{w}(i)}{\sum_i \bar{w}(k)} S (\bar{w}(i) / \hat{y}(i); i, q) \sum_i \bar{w}(k) \\ &= \bar{S} (p^h + p^e + E), \end{aligned}$$

where \bar{S} is the initial-wealth-weighted average savings rate and E is aggregate initial wealth not invested in long-lived assets, as defined in (6).

Suppose for the moment that new asset supply is the same for equity and houses, that is, $f^e = f^h = f$. The intuition generated by this case is helpful for interpreting our results more generally. The aggregate value of these assets now becomes

$$(8) \quad p^h + p^e = \frac{\bar{S}E - D}{1 + f - \bar{S}}.$$

Other things – including expectations – equal, the value of long-lived assets decreases when the ROE supplies more of these assets. It also decreases when the ROE raises more funds in the bond market. Here it matters that our calibration strategy takes the market value of bond supply as given (rather than, say, the face value). Higher bond supply in our exercises always subtracts a fixed amount from the demand for other savings vehicles, such as houses and equity and thus lowers the value of those vehicles. Finally, the value of long-lived assets increases with any event that increased the average savings rate in the household sector. In particular, it depends on household expectations only through their effect on the average savings rate.

⁹Asset demand depends on the prices of long-lived assets only through the effects of those prices on initial wealth. The reason is that return expectations are exogenous, except for bonds where they depend on the market interest rate and expected inflation.

An important implication of equation (8) is that the *total* value of long-lived assets is determined by consumption-savings considerations, not by portfolio choice considerations. Indeed, at the household level, the optimal savings rate is chosen to smooth consumption over time, given the total return on wealth. In contrast, the weights α on the individual assets are chosen to generate the highest possible total return on wealth, trading off risk and return. The total value of long-lived assets depends strongly on aspects of the household problem that affect the savings decision, such as the discount factor, the intertemporal elasticity of substitution, the age distribution and expectations of labor income growth. It is less sensitive to factors that matter for portfolio choice, such as the relative returns on different assets. The latter factors are instead important for the individual asset prices, as well as the bond price, which must adjust to ensure that all three equations in (7) hold.

Aggregation

The distribution of household characteristics – in particular age and endowments – affects asset prices via the average savings rate (cf. (8), for example) as well as via average portfolio weights. Since households have homothetic preferences, it is not obvious that features of the distribution other than the means matter for prices.¹⁰ Recent work on calibrated incomplete markets models has derived “approximate aggregation” results: moments of the wealth distribution beyond the mean often have little effect on equilibrium asset prices (see Krusell and Smith 2006 for a survey).

Approximate aggregation obtains if individual savings $S(\bar{w}(i)/\hat{y}(i); i, q)$ \bar{w} , viewed as a function of initial wealth \bar{w} , are well approximated by affine functions with a common slope:

$$(9) \quad S(\bar{w}(i)/\hat{y}(i); i, q) \bar{w} \approx a(i, q, \hat{y}(i)) + b(q) \bar{w},$$

at least on the support of the wealth distribution. Here the key property is that b does not depend on any individual characteristics i . Assuming equal supply of equity and houses as in (8): the total value of long lived assets is then

$$(10) \quad p^h + p^e = \frac{\bar{a} + bE - D}{1 + f - b},$$

where \bar{a} is the population-weighted average of the intercepts a . The wealth-GDP ratio thus depends on the distribution of income only through the aggregate E .

The top panel of Figure 7 plots the ratio of savings – or terminal wealth – to the permanent component of income \hat{y} as a function of the wealth-to-income ratio \bar{w}/\hat{y} , for the same age groups as in Figure 5. Holding age fixed, savings are well approximated by an affine function in \bar{w} , if \bar{w} is high relative to \hat{y} .¹¹ Moreover, the intercept is negative and proportional to \hat{y} : rich households with a higher permanent component of income save less overall, but they save the same fraction of every additional

¹⁰In a standard general equilibrium model with identical homothetic preferences, the wealth distribution is irrelevant for prices: homotheticity leads to demand functions that are linear in wealth; if preferences are also identical, the slopes are all the same and aggregate demand is linear in aggregate wealth.

¹¹Comparison of Figures 5 and 7 shows that equal savings rates across individuals are sufficient, but not necessary to obtain common slopes b of the savings functions (9).

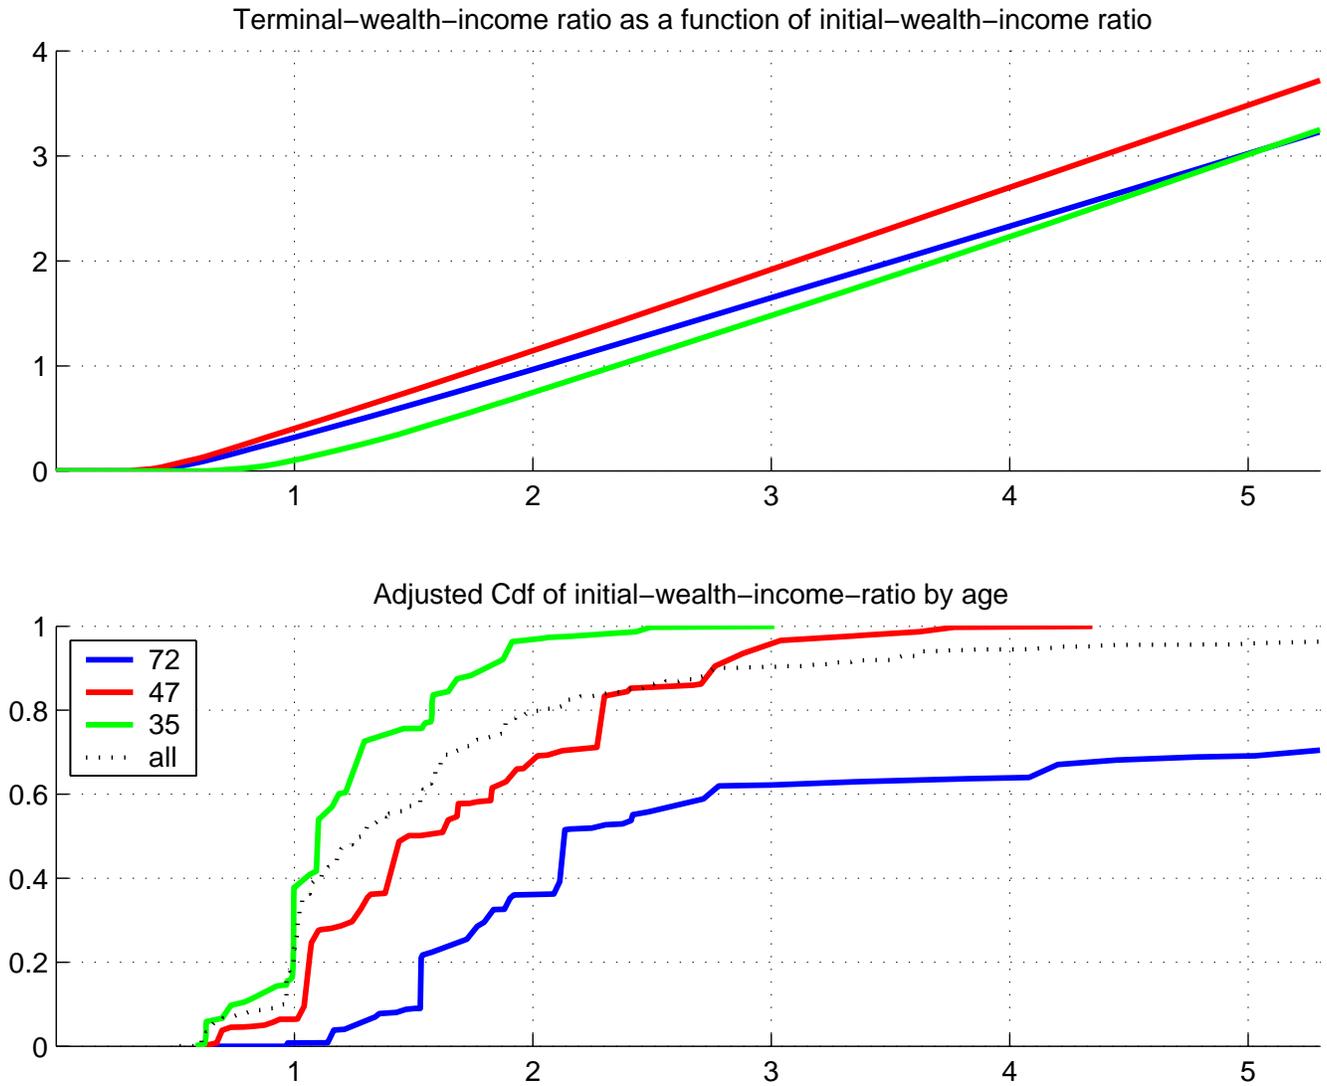


Figure 7: Aggregation Results

unit of wealth. While approximate aggregation obtains for rich households within an age group, the savings function is not affine when \bar{w} is low relative to \hat{y} . This is true especially for the youngest agents, for whom the nonnegativity constraint on savings is most relevant. In addition, the slopes differ by age, which suggests that the wealth distribution between cohorts matters for prices.

Our framework allows to assess how much heterogeneity in wealth matters for prices by comparing the measured wealth distribution to the model implied savings function. The bottom panel of Figure 7 plots a cumulative distribution functions for the wealth-to-income ratio \bar{w}/\hat{y} by age. For easier interpretation, the cdf was adjusted so that integrating the savings functions for some age group in the top panel against the density derived from the corresponding cdf in the bottom panel yields aggregate

savings (up to a scaling factor).¹² Approximate aggregation within an age group thus obtains if most of the mass is in a region where the savings function is linear. While this is true for old and middle aged households who tend to have a lot of wealth relative to income, the majority of young agents (29-35) is in a region where the savings function is still convex.

VI Supply, Demographics and Asset Prices 1968-95

Panels A and B of Table 4 compare the baseline model to the data not only for 1995, but also for 1968 and 1978. The difference between 1968 and 1995 is that the wealth-GDP ratio as well as the nominal interest rate were both slightly lower in 1968. The baseline model accounts for most of this difference because the variance of income shocks was lower in 1968 – lower precautionary savings lowered wealth and interest rates. With a spread of 2% between borrowing and lending rates, the model matches gross borrowing and lending in 1968. A fairly small drop in the spread – 40 basis points – is thus sufficient to account for the increase in the volume of credit.

TABLE 4: BASELINE RESULTS

year	beliefs	wealth/ GDP	portfolio weights			lend./ GDP	borr/ GDP	PD ratios		interest rate nominal
			bonds	housing	stocks			housing	stocks	
PANEL A: DATA										
1968		2.41	0.17	0.51	0.32	0.61	0.20	20.1	32.3	0.057
1978		2.08	0.16	0.68	0.16	0.56	0.23	24.4	18.6	0.084
1995		2.51	0.15	0.59	0.26	0.70	0.31	20.4	23.9	0.061
PANEL B: MODEL										
1968	Baseline	2.46	0.17	0.58	0.25	0.61	0.2	23.8	25.1	0.062
1978	Baseline	2.09	0.16	0.61	0.23	0.54	0.20	22.3	25.9	0.064
1995	Baseline	2.52	0.15	0.60	0.25	0.70	0.31	20.9	23.2	0.062

NOTE: Panel A reports the aggregate portfolio weights on bonds, housing and stocks from Figure 2; the gross borrowing and lending numbers from Section the SCF of 1962, the constructed numbers for 1978 and the SCF of 1995; the wealth-to-GDP ratio from Figure 1; the price-dividend ratios for housing and stocks together with the nominal 6-year interest rate. Panel B reports the results computed from the model with baseline beliefs.

Household portfolios in 1978 were very different from those in 1968 or 1995: wealth as a percent of GDP was much smaller, and there was a strong portfolio shift from stocks into houses. The model with baseline expectations held fixed delivers the first fact, but not the second. The wealth to GDP ratio drops to about twice GDP in both the model and the data. However, the portfolio allocation in the

¹²More specifically, each households' weight was multiplied by the ratio of its individual $\hat{y}(i)$ to the aggregate $\hat{y}(i)$.

model remains essentially the same as in 1995. As a result, the price dividend ratios of houses falls and that of houses rises, in contrast to what happened in the data.

Three changes in fundamentals are important for the drop in the wealth-to-GDP ratio in 1978. The first is the change in the distribution of endowments illustrated in Figure 4. The special feature of the 1978 endowment distribution is that a larger fraction of the funds available for investment resides with the very youngest and oldest cohorts. As shown in the last section, the model predicts that these cohorts have small savings rates, which leads to lower wealth-to-income ratios and pushes interest rates up. This effect is reinforced by the effects of low ex post real interest rates, which also reduce the ratio of initial wealth to income and hence savings. A counteracting force is the reduction in bond supply documented in Figure 3 which tends to lower interest rates and slightly raises the wealth-GDP ratio. Taken together, these effects produce relatively stable interest rates and a low wealth-GDP ratio.

TABLE 5: ADDITIONAL EXPERIMENTS FOR 1968 AND 1978

year	experiment	wealth/ GDP	portfolio weights			lend./ GDP	borr/ GDP	PD ratios		int. rate nom.
			bonds	housing	stocks			housing	stocks	
1968	(1) Baseline	2.46	0.17	0.58	0.25	0.61	0.20	23.8	25.1	0.062
1968	(2) 1978 Bond Supply	2.55	0.13	0.60	0.27	0.58	0.25	25.2	28.4	0.060
1968	(3) 1978 Bond Endow.	2.34	0.17	0.59	0.24	0.59	0.18	22.6	23.2	0.063
1978	(4) Baseline	2.09	0.16	0.61	0.23	0.54	0.20	22.3	25.9	0.064
1978	(5) 1968 Bond Market	2.06	0.20	0.60	0.20	0.58	0.17	21.4	22.9	0.066

NOTE: The table reports model results for various years and counterfactuals.

Table 5 reports experiments that isolate changes in each of the three factors, leaving the others fixed at their 1968 values. The experiment in row 2 considers bond supply: we retain the 1968 distribution of endowments and change only the supply of new bonds to the 1978 value, a reduction of 25%. We find that house and stock prices increase enough to increase the wealth-GDP values while the interest rate falls. Intuitively, a reduction in the supply of any savings vehicle – here bonds – raises the prices of all savings vehicles; it thus raises the wealth to GDP ratio and lowers the nominal interest rate.

Lower nominal interest rates in turn lead to more borrowing within the household sector – higher supply of bonds from borrowing households thus partly offsets the shortfall of bond supply from the rest of the rest of the economy. Because of the collateral constraint, borrowing goes along with more investment in housing relative to stocks – house prices increase proportionately more than stock prices and there is a portfolio shift from bonds into housing.

The experiment in row 3 isolates the role of surprise inflation on bond endowments. It retains the 1968 bond supply as well as the 1968 distribution of income, house and stock endowments. However, it constructs bond endowments by updating bond holdings with an interest rate factor that is about 3% lower than the factor used in the 1968 baseline case. The new factor thus lowers the initial wealth

of lenders and increases the initial wealth of borrowers; its value is selected to make aggregate payoffs on bonds as a percent of GDP as low as in 1978, where inflation significantly reduced the ex post real interest rate.

Surprise inflation has the opposite effects from a reduction in asset supply: a lower bond endowment lowers initial wealth and thus reduces the demand for all savings vehicles. As a result, all asset prices fall, as does the wealth-GDP ratio. The reason is that the household sector on aggregate is a net lender so that for the majority of households (in wealth-weighted terms) the wealth-income ratio and the savings rate go down. Of course, at the same time, borrower households experience an increase in their wealth-income ratio, so that the distribution of wealth-income ratios becomes less dispersed. This explains the drop in gross credit volumes from the experiment.

Finally, row 5 uses the 1978 distribution, including bond endowments, but fixes the bond supply at its 1968 value. In addition, it increases the bond endowments by raising the interest rate factor by about 3%, so that the aggregate bond endowments is also at the higher 1968 value. This is a way to isolate the effect of the income and endowment distribution from the other two factors. The main result is a drop in the wealth-GDP ratio that is essentially the same as in the data. It is driven by the lower savings rates of the 1978 population.

VII The Effect of Beliefs about Inflation

In this section, we consider the effects of inflation on the price-dividend ratios of real assets in the 1970s. We use our model to quantify the extent to which changes in both expected inflation and inflation uncertainty contributed to higher house prices and lower stock prices. Results are reported in Table 6. Rows 1 and 2 simply repeat the statistics from Table 4 for the data and model results based on baseline beliefs, respectively. In this section, we show how model results change as we vary beliefs about inflation.

We begin by examining the first moment, that is, expected inflation. To measure inflation expectations in the 1970s, we turn to the Michigan Survey of Consumers. Figure 8 plots population median forecasts since 1979 together with median forecasts for different cohorts in the survey. The figure illustrates two key stylized facts. First, median inflation expectations across all households are high during the late 1970s and early 1980s, and then come down. Second, households disagree systematically by age. During the Great Inflation, older households expected much lower inflation than younger households; the discrepancy vanished as inflation subsided in the 1990s.

The exercise in row 3 of Table 6, labeled “Survey Inflation Exp ”, assumes that the inflation rate expected by an age cohort in the model is equal to the corresponding cohort median in the Michigan Survey. The experiment generates a dramatic shift in portfolios towards housing, driven by an increase in house prices and a drop in stock prices. At the same time, the nominal interest rate rises sharply to over 10% and there is a credit boom: gross borrowing and lending increases by close to 10 percentage points.

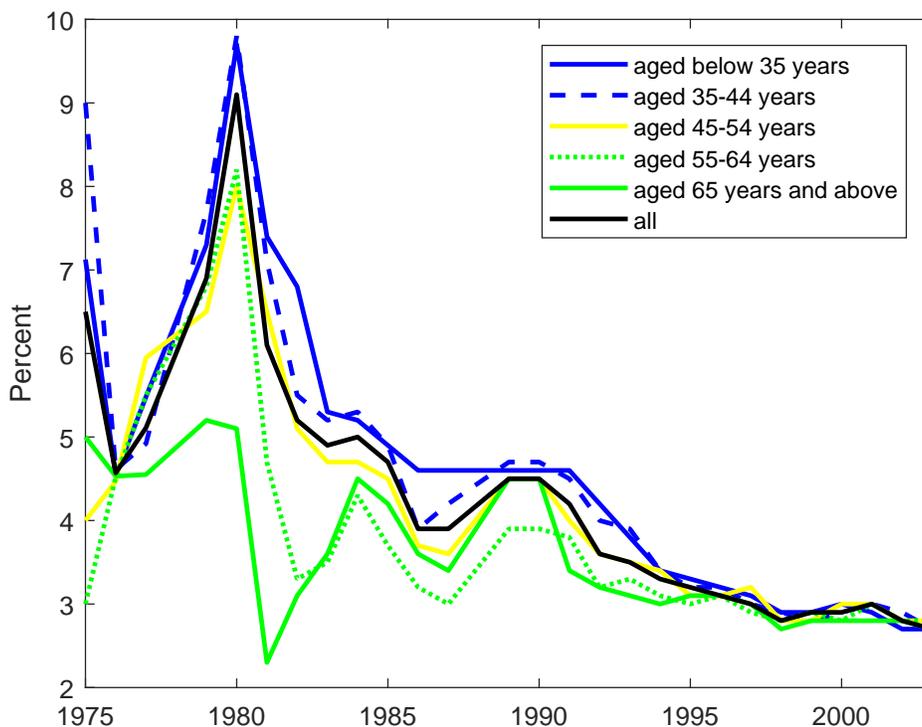


Figure 8: Median inflation expectations over the next 5 years, Michigan Survey of Consumers. Each line represents inflation expectations for households in the age group indicated in the legend. The black solid line is the median expectation across all households.

Average Expected Inflation and Taxes

To understand how these results come about, it is helpful to separate higher *average* inflation from disagreement. Row 3, labeled “High Infl Exp” focuses on the former: it increases expected inflation for all households from the baseline number of 4% to the median 5-year inflation forecast from the Michigan survey, averaged over the trading period, which is 6.3%. We find essentially the same aggregate portfolio adjustment and hence house and stock prices as with disagreement in Row 2. However, the credit boom is now gone, and the nominal interest rate is about one percentage point lower.

With higher average expected inflation in Row 3, all effects are due to nominal rigidities in the tax system: taxation of nominal returns makes housing more attractive as an investment relative to stocks and bonds. To see this, let r_t^e denote the pretax real return on equity and let π_t denote the inflation rate. Under our assumption of a single effective tax rate on nominal equity returns, the after-tax real return on equity is $(1 - \tau) r_t^e - \tau \pi_t$. The expected after-tax real return on equity is thus decreasing in expected inflation. In contrast, we have assumed that housing returns are not taxed, so that the expected real return on housing does not depend on expected inflation. As expected inflation increases, real estate thus become more attractive relative to equity.

The return on nominal bonds is $(1 - \tau) i_t - \pi_t = (1 - \tau) r_t^b - \tau \pi_t$, where i_t is the nominal interest rate and $r_t^b = i_t - \pi_t$ is the pretax real return. Given a nominal rate, holding bonds thus also becomes

less attractive as expected inflation increases, whereas mortgage borrowing becomes more attractive because interest is tax deductible. Put differently, the taxation of nominal interest induces an inflation premium on top of the increase in the nominal rate warranted by the Fisher equation. This is why the nominal interest rate increases by more than the increase in expected inflation.

Disagreement about Expected Inflation

Heterogeneity in inflation expectations matters in our model because it introduces disagreement about real returns that correlates with the propensity to take risks. Two effects explain the differences between Rows 3 and 4 of Table 6. First, old households who expect lower inflation perceive relatively higher after-tax returns on stocks. As a result, the shift to housing discussed above is muted for old households who expect lower inflation than the median household. Since old households are also more important holders of equity on aggregate, the disagreement case in Row 3 shows a higher weight on stocks as well as a higher stock price.

The second effect is that disagreement about real interest rates generates gains from trade in the credit market. Young households – who expect high inflation – perceive a lower real interest rate than old households who expect low inflation. The result is a credit boom: young households who are natural borrowers believe that borrowing is a bargain and borrow from old households, who are happy to lend. Given our distribution of endowments, the relative attraction of borrowing is so large that the interest rate has to rise substantially to choke off the credit boom. Disagreement thus moves the real rate, holding fixed average expected inflation.

Quantitatively, the changes in household portfolios in row 3 are more than enough to match the portfolio shift of the 1970s shown in row 1. However, the nominal interest rate of 10.1% is significantly higher than the observed rate of 8.4% and the credit boom generated by disagreement is too large. This reveals a basic tension between using point estimates of expected inflation and fitting moments from the credit market. The result is perhaps not surprising. Indeed, the Great Inflation increased the volatility of inflation, but only its mean, and it is likely that households took this into account. Moreover, there is considerable measurement error in 1970s survey data. We now attempt to address both issues.

Inflation Uncertainty

Given the increase in many measures of inflation uncertainty during the 1970s, it is perhaps not surprising that an asset pricing model that varies only the first moments of beliefs cannot fit the data. We now turn to the quantitative effects of inflation volatility. Unfortunately, we do not have data on perceived uncertainty directly from survey. Our strategy is thus to back out beliefs and assess whether this provides a parsimonious account of savings, portfolio choice and asset prices.

To provide some intuition for how inflation uncertainty works in our model, we perform a counterfactual that starts from the baseline and simply doubles inflation volatility from 3% to 6% p.a. Results are reported in Row 5 of Table 6, labeled “High Infl Vol”. The main difference to the baseline results is that there is now less nominal borrowing and lending within the household sector and a higher nominal interest rate. Intuitively, higher conditional volatility of inflation increases the conditional volatility

of real bond returns and the ex post real cost of mortgage financing. Nominal instruments thus become less attractive to both borrowers and lenders. The nominal interest rate increases because the household sector as a whole is a net lender and must be compensated if it is to hold the nominal debt of the rest of the economy.

Increasing inflation volatility has little effect on the overall wealth-GDP ratio. Indeed, it is a comparative static that affects the returns on one asset, leaving those on other assets unchanged. The overall return on wealth and the savings decision are not affected much and the equilibrium wealth-GDP ratio stays essentially the same. Aggregate portfolio weights also remain unchanged. On the one hand, households are required to hold the same value of bonds issued by the rest of the economy in equilibrium, so that the portfolio weight on bonds remains unchanged if the wealth-GDP does not change. On the other hand, the effect of a change in the relative volatility of after tax returns on stocks and houses is quantitatively small so that the weights on stocks and houses also move little.

TABLE 6: THE EFFECTS OF INFLATION

beliefs	wealth/ GDP	portfolio weights			lend./ GDP	borr./ GDP	PD ratios		nominal int. rate
		bonds	housing	stocks			housing	stocks	
(1) — 1978 Data —	2.08	0.16	0.68	0.16	0.56	0.23	24.4	18.6	0.084
(2) Baseline	2.09	0.16	0.61	0.23	0.54	0.20	22.3	25.9	0.064
(3) Survey Infl Exp	2.02	0.16	0.70	0.14	0.69	0.36	24.5	15.1	0.102
(4) High Infl Exp	2.01	0.17	0.71	0.12	0.52	0.18	24.9	13.3	0.091
(5) High Infl Vol	2.08	0.16	0.62	0.22	0.50	0.17	22.4	25.1	0.066
(6) Survey Infl Exp + Hi Infl Vol	2.02	0.16	0.70	0.14	0.56	0.23	24.5	15.1	0.108
(7) Adj Infl Exp + Hi Infl Vol	2.03	0.16	0.69	0.15	0.55	0.22	24.2	16.5	0.085

NOTE: Rows 1 repeats the 1978 data, while row 2 contains results with baseline expectations from Table 3. The other scenarios change expectations over the next six years away from the baseline. Row 3 (“Survey Infl Exp”) sets every age cohort’s expected inflation in the model equal to the cohort median in the Michigan Survey, averaged over 1975-80. Row 4 (“High Infl Exp”) increases households’ expected inflation from the baseline 4% to the median 6.3% in the Michigan Survey of Consumers, averaged over 1975-1980. Row 5 (“High Infl Vol”) doubles the conditional volatility of inflation. Row 6 (“Survey Infl Exp + Hi Infl Vol”) uses expected cohort inflation rates (as in row 3) and increase the volatility of inflation by a factor of 3.5 Row 7 (“Adj Infl Exp + Hi Infl Vol”) shifts expected inflation down by 1pp per year for all cohorts.

We now ask whether there is a parsimonious parametrization of beliefs that avoids the tension between effects on long-lived asset markets and the credit market. We proceed in two steps. First, we ask what the perceived volatility of inflation would have to be so the model matches the volume of credit. Row 6 of Table 5 starts from heterogeneous expectations (row 3), but increases the volatility of inflation by a factor of 3.5. As expected from the comparison of rows 1 and 5 above, the only effect is that credit drops to match the data, and there is a small increase in the interest rate. We conclude

that there are no important interaction between disagreement and perceived volatility – the two belief parameters impact different moments of the data.

The final experiment in Row 7 shifts down the schedule of expected inflation by 1pp per year for all cohort. We thus continue to match cross-cohort differences in expected inflation and only adjust the mean. The adjustment shift is within two standard errors of the point estimates for the cohort medians, and hence still represents a plausible scenario, given the uncertainty about survey expectations during this period. The result is that we can now jointly match wealth-GDP, portfolio weights, the volume of credit and interest rates. We conclude that introducing inflation expectations that are broadly in line with survey evidence can account for the strong portfolio shift across assets. Together with demographics – which according to the last row in Table 5 is responsible for joint downward movement in long-lived wealth positions – we arrive at an account for the 1970s driven primarily by the two unique events of the decade.

References

- Abel, Andrew B.**, "The Effects of a Baby Boom on Stock Price and Capital Accumulation in the Presence of Social Security," *Econometrica*, 2003, 71 (2), 551–578.
- Ameriks, John, Joseph Briggs, Andrew Caplin, Matthew D. Shapiro, and Chris Tonetti**, "Long-Term-Care Utility and Late-in-Life Saving," *Forthcoming Journal of Political Economy*, 2020.
- Campbell, John Y. and Joao F. Cocco**, "Household Risk Management and Optimal Mortgage Choice.," 2003.
- Chien, YiLi, Harold Cole, and Hanno Lustig**, "Is the Volatility of the Market Price of Risk Due to Intermittent Portfolio Rebalancing?," *American Economic Review*, 2012, 102 (6), 2859–2896.
- , —, and —, "Implications of heterogeneity in preferences, beliefs and asset trading technologies in an endowment economy," *Review of Economic Dynamics*, 2016, 20, 215–239.
- Cocco, Joao F.**, "Portfolio Choice in the Presence of Housing," *Review of Financial Studies*, 2005, 18, 535–567.
- Cogley, Timothy and Thomas J. Sargent**, "Evolving Post-World War II U.S. Inflation Dynamics," *NBER Macroeconomics Annual*, 2001, 16, 331–373.
- Constantinides, George M. and Darrell Duffie**, "Asset Pricing with Heterogeneous Consumers.," *Journal of Political Economy*, 1996, 104 (2), 219–40.
- , **John B. Donaldson, and Rajnish Mehra**, "Junior Can't Borrow: A New Perspective on the Equity Premium Puzzle," *Quarterly Journal of Economics*, 2002, 117, 269–296.
- DeNardi, Mariacristina, Eric French, and John B. Jones**, "Why do the Elderly Save? The Role of Medical Expenses," *Journal of Political Economy*, 2010, 118 (1), 39–75.
- , —, and **John Jones**, "Savings After Retirement: A Survey," *The Annual Review of Economics*, 2016, 8, 177–204.
- Doepke, Matthias and Martin Schneider**, "Inflation and the Redistribution of Nominal Wealth," *Journal of Political Economy*, 2006, 114 (6), 1069–1097.
- Eggertsson, Gauti B., Neil R. Mehrotra, and Jacob A. Robbins**, "A Model of Secular Stagnation: Theory and Quantitative Evaluation," *American Economic Journal: Macroeconomics*, 2019, 11 (1), 1–48.
- Favilukis, Jack, Sydney C. Ludvigson, and Stijn van Nieuwerburgh**, "The Macroeconomic Effects of Housing Wealth, Housing Finance, and Limited Risk Sharing in General Equilibrium.," *Journal of Political Economy*, 2017, 125 (1), 140–223.
- Feldstein, Martin**, "Inflation and the Stock Market.," *American Economic Review*, 1980, 70 (5), 839–847.
- Fernandez-Villaverde, Jesus and Dirk Krueger**, "Consumption and Saving over the Life Cycle: How Important are Consumer Durables?," *Macroeconomic Dynamics*, 2011, 15 (5), 725–770.
- Flavin, Marjorie and Takashi Yamashita**, "Owner-Occupied Housing and the Composition of the Household Portfolio.," *American Economic Review*, 2002, pp. 345–62.
- Gagnon, Etienne, Benjamin K. Johansson, and David Lopez-Salido**, "Understanding the New Normal: The Role of Demographics," *Finance and Economics Discussion Series of the Board of Governors*, 2016, 2016-080.
- Geanakoplos, John, Michael Magill, and Martine Quinzii**, "Demography and the Long-run Predictability of the Stock Market.," *Brookings Papers on Economic Activity*, 2004, 1, 241–307.
- Gomes, Francisco and Alex Michaelides**, "Optimal Life-Cycle Asset Allocation: Understanding the Empirical Evidence," *Journal of Finance*, 2005, 60 (2), 869–904.

- Gourinchas, Pierre-Olivier and Jonathan Parker**, "Consumption over the Life Cycle," *Econometrica*, 2002, 70, 47–89.
- Hansen, Lars Peter and Kenneth J. Singleton**, "Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models," *Econometrica*, 1982, 50 (5), 1269–1286.
- Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante**, "The Macroeconomic Implications of Rising Wage Inequality in the United States," *Journal of Political Economy*, 2010, 118 (4), 681–722.
- Heaton, John and Deborah J. Lucas**, "Evaluating the Effects of Incomplete Markets on Risk Sharing and Asset Pricing," *Journal of Political Economy*, 1996, 104, 443–487.
- and —, "Portfolio Choice and Asset Prices: The Importance of Entrepreneurial Risk," *Journal of Finance*, 2000, 55, 1163–1198.
- Ilut, Cosmin and Martin Schneider**, "Ambiguous Business Cycles," *American Economic Review*, 2014, 104, 2368–2399.
- Jagannathan, Ravi and Narayana Kocherlakota**, "Why Should Older People Invest Less in Stocks Than Younger People?," *Quarterly Review of the Federal Reserve Bank of Minneapolis*, 1996, 20 (3), 11–23.
- Justiniano, Alejandro, Giorgio E. Primiceri, and Andrea Tambalotti**, "Credit Supply and the Housing Boom," *Journal of Political Economy*, 2019, 127 (3), 1317–1350.
- Krueger, Dirk and Felix Kubler**, "Computing equilibrium in OLG models with stochastic production," *Journal of Economic Dynamics and Control*, 2004, 58 (7), 1411–1436.
- Krusell, Per and Anthony Smith**, "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy*, 1998, 106 (5), 867–896.
- Landvoigt, Tim, Monika Piazzesi, and Martin Schneider**, "The Housing Market(s) of San Diego," *American Economic Review*, 2015, 105 (4), 1371–1407.
- Malmendier, Ulrike and Stefan Nagel**, "Learning from Inflation Experiences," *Quarterly Journal of Economics*, 2016, 131 (1), 53–87.
- Mankiw, Gregory N. and David Weil**, "The Baby Boom, the Baby Bust and the Housing Market," *Regional Science and Urban Economics*, 1989, 19, 235–258.
- Ortalo-Magne, Francois and Sven Rady**, "Housing Market Dynamics: On the Contribution of Income Shocks and Credit Constraints.," *Review of Economic Studies*, 2006, 73 (2), 459–485.
- Piazzesi, Monika and Martin Schneider**, "Housing and Macroeconomics," *Handbook of Macroeconomics*, 2016, 2, 1547–1640.
- , —, and **Selale Tuzel**, "Housing, Consumption and Asset Pricing," *Journal of Financial Economics*, 2007, 83, 531–569.
- Poterba, James**, "House Price Dynamics: The Role of Tax Policy and Demographics.," *Brookings Papers on Economic Activity*, 1991, 2, 143–203.
- Sialm, Clemens**, "Tax Changes and Asset Pricing," *American Economic Review*, 2009, 99 (4), 1356–1383.
- Storesletten, Kjetil, Chris I. Telmer, and Amir Yaron**, "Asset pricing with idiosyncratic risk and overlapping generations," *Review of Economic Dynamics*, 2007, 10 (4), 519–548.
- Summers, Lawrence H.**, "Inflation, the Stock Market, and Owner-Occupied Housing," *American Economic Review*, 1981, 71, 429–434.

Yao, Rui and Harold Zhang, "Optimal Portfolio Choice with Risky Housing and Borrowing Constraints.," *Review of Financial Studies*, 2005, 18, 197–239.

Zeldes, Stephen P., "Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence," *Quarterly Journal of Economics*, 1989, 104, 275–298.

Appendix

A Assets and Income

This appendix describes how we map observed asset classes into the three assets present in the model, and how we measure aggregate and household level asset holdings, as well as non-asset income.

A. Data and Definitions

We map the three assets in the model to three broad asset classes in US aggregate and household statistics. Our main data sources are the Financial Accounts of the United States (FAUS) for aggregates and the Survey of Consumer Finances (SCF) for individual positions. To make these data sets comparable, we must ensure that aggregates match. As shown by Antoniewicz (2004), the match is good for most asset classes in both 1989 and 1995, after a few adjustments. However, our own computations show that the match for nominal assets is bad for the 1962 SCF. For some classes of assets, especially short-term deposits, the SCF aggregates are only about 50% of the FAUS aggregates. Apparent underreporting of short-term nominal assets is also present in later SCFs, but is less severe. To achieve a comparable time series of positions, we assume throughout that the FAUS aggregates are correct and that individual positions in the SCF suffer from proportional measurement error. We then multiply each individual position by the ratio of the FAUS aggregate and the SCF aggregate for the same asset class.

Asset classes

We identify *equity* with shares in corporations held and controlled by households, including both publicly traded and closely held shares, and both foreign and domestic equity. We also include shares held indirectly through investment intermediaries if the household can be assumed to control the asset allocation (into our broad asset classes). We take this to be true for mutual funds and defined contribution (DC) pensions plans. For these intermediaries, while the fund manager determines the precise composition of the portfolio, the household typically makes the decision about equity versus bonds by selecting the type of fund.

We thus consolidate mutual funds and DC pension funds. For example, when households own a mutual fund, an estimate of the part of the fund invested in stocks is added to stock holdings. In contrast, we do not include equity held in defined benefit (DB) pension plans, since the portfolios of these plans are not controlled by households themselves. Instead, DB plans are treated as a tax-transfer system sponsored by the rest of the economy (in practice, the corporate sector or the government). We also do not include noncorporate business, which is treated partly as real estate and partly as labor income, as described below.

We construct an annual series for the aggregate value of household sector equity holdings. Our starting point is the published series in the FAUS. We cannot use that series directly, since it contains (i) the market value of the equity component of foreign direct investment (that is, equity positions by foreigners in excess of 10% of shares in a US corporation) and (ii) the market value of equity held by DB pension funds. We estimate the equity component of FDI using data on the International Investment

Position from the Bureau of Economic Analysis. Shares held by defined benefit pension funds are available from the FAUS. Our series is obtained by subtracting (i) and (ii) from the FAUS series on household equity holdings.

We derive estimates of net new equity purchased by households using a similar correction of FAUS numbers. Finally, our concept of dividends equals dividends received by households from the National Income and Product Accounts (NIPA) less dividends on their holdings in DB pension plans. We use the numbers on value, dividends and new issues to calculate price dividend ratios and holding returns on equity. The properties of the return series are discussed in Appendix B below. For household-level positions, we use the Survey of Consumer Finances, which also contains direct holdings of publicly-traded and closely-held shares, as well as an estimate of equity held indirectly through investment intermediaries.

Our concept of *residential real estate* contains owner-occupied housing, directly held residential investment real estate, as well as residential real estate recorded in the FAUS/NIPA as held indirectly by households through noncorporate businesses. This concept contains almost all residential real estate holdings, since very few residential properties are owned by corporations. To construct holdings of tenant-occupied residential real-estate at the individual level, we start from the SCF numbers and then add a proportional share of the household's noncorporate business position. This share is selected so that our aggregate of tenant-occupied real estate over all households matches the corresponding value from the FAUS. We take housing dividends to be housing consumption net of maintenance and property tax from NIPA. For net purchases of new houses, we use aggregate residential investment from NIPA. As with equity, the annual series for holdings, dividends and new issues give rise to a return series, discussed below.

Household *bond* holdings are set equal to the net nominal position, that is, the market value of all nominal assets minus the market value of nominal liabilities. Nominal assets again include indirect holdings through investment intermediaries. To calculate market value, we use the market value adjustment factors for nominal positions in the U.S. from Doepke and Schneider (2006). In line with our treatment of tenant-occupied real estate, we assign residential mortgages issued by noncorporate businesses directly to households. At the individual level, we assign a household mortgages in proportion to his noncorporate business position, again with a share selected to match the aggregate value of residential mortgages from the FAUS.

Non-Asset Income

Our concept of non-asset *income* comprises all income that is available for consumption or investment, but not received from payoffs of one of our three assets. We start by constructing an aggregate measure of such income from NIPA. Of the various components of worker compensation, we include only wages and salaries, as well as employer contributions to DC pension plans. We do not include employer contributions to DB pension plans or health insurance, since these funds are not available for consumption or investment. However, we do include benefits disbursed from DB plans and health plans. Also included are transfers from the government. Finally, we subtract personal income tax on

non-asset income.

Non-asset income also includes dividends from noncorporate business except those attributable to residential real estate. To construct the latter concept of noncorporate dividends, we use the aggregate price-dividend ratio of housing to estimate the housing dividend provided by an individual's private business. Our approach essentially splits up the noncorporate sector into a real estate component that is very capital intensive and relies heavily on debt finance, and a rest that is much more labor intensive. Indeed, the capital stock of the noncorporate nonfinancial sector in 1995 was \$3.6trn, of which \$2.4trn was residential real estate.

Given the aggregate series of income, we apply the same conventions to individual income in the SCF to the extent possible. A problem is that the SCF does not report employer contributions to DC plans and only reports pretax income for all items. To address this issue, we apply a proportional tax rate to pretax non-asset income reported in the SCF, where the tax rate is chosen such that aggregate non-asset income is equal to its counterpart from NIPA. The outcome is an income distribution that matches with NIPA at the aggregate level.

B. Measuring the Distribution of Endowments

Consumers in our model are endowed with both assets and non-asset income. To capture decisions made by the cross-section of households, we thus have to initialize the model for every period t with a *joint* distribution of asset endowment and income. We derive this distribution from asset holdings and income observed in the previous period $t - 1$. Limits on data availability imply that we have to resort to different approaches for the different years. For 1995, the data situation is best, since we can use the SCF in the 4th year of period t together with the SCF from the 4th year of period $t - 1$. We describe our strategy first for this case. We then explain how it is modified for earlier years where less data are available.

Approximating the distribution of households

In principle, one could use all the households in SCF and update them individually. This would lead to a large number of agents and consequently a large number of portfolio problems would have to be solved. We simplify by approximating the distribution of endowments and income with a small number of household types. First, we sort households into the same nine age groups described at the beginning of this section. Within each age group, we then select 6 subclasses of SCF households. We start by extracting the top 10% of households by net worth. Among the bottom 90% net worth, we divide by homeowner and renter. We then divide homeowners further into "borrowers" and "lenders." Here a household is a borrower if his net nominal position – nominal asset minus nominal liabilities – is negative. We further subdivide each homeowner category into high/low wealth-to-income ratios.

The above procedure splits up households into $9 \times 6 = 54$ different cells. We assume that all households that fall into the same cell are identical and compute asset positions at the cell level. The SCF survey weights determine the cell population. Naturally, the procedure loses some features of the true distribution due to aggregation. However, it ensures that key properties of the distribution

that we are interested in are retained. In particular, because very few among the top 10% are net nominal borrowers, gross borrowing and lending are very close in the true and the approximating distributions. In addition, the approximating distributions retains asset positions conditional on age as well as conditional on wealth when net worth is split as top 10% and the rest, our key measures of concentration described.

Asset endowments for a transiting individual household

Consider the transition of an individual household's asset position from period $t - 1$ into period t . We have treated both stocks and houses as long-lived trees and we normalize the number of trees carried into the period by consumers to one. We can thus measure the household's endowment of a long-lived asset from its share in total market capitalization of the asset in period $t - 1$. The SCF does not contain consumption data. Using the language introduced in the discussion of the budget constraint (4), we can thus measure either initial or terminal wealth in a given period, but not both. We assume that terminal wealth can be directly taken from the survey. The initial supply of assets is normalized to one, so that the initial holdings of housing $\bar{\theta}_t^h$ and stocks $\bar{\theta}_t^e$ are the agent's *market shares* in period $t - 1$. For each long-lived asset $a = h, e$, suppose that $w_{t-1}^a(i)$ is the market value of investor i 's position in $t - 1$ in asset a . Now we can measure household i 's initial holdings as

$$\begin{aligned}\bar{\theta}_t^a(i) &= \theta_{t-1}^a(i) = \frac{w_{t-1}^a(i)}{\sum_i w_{t-1}^a(i)} = \frac{p_{t-1}^a \theta_{t-1}^a(i)}{p_{t-1}^a \sum_i \theta_{t-1}^a(i)} \\ &= \text{market share of household } i \text{ in period } t - 1.\end{aligned}$$

Updating Nominal Positions

For the nominal assets, the above approach does not work since these assets are short-term in our model. Instead, we determine the market value of nominal positions in period $t - 1$ and update it to period t by multiplying it with a nominal interest rate factor. In particular, suppose that $w_{t-1}^b(i)$ is the market value of investor i 's net nominal positions in $t - 1$ and that

$$\begin{aligned}\bar{\theta}_{t-1}^b(i) &= \frac{w_{t-1}^b(i)}{\sum_i w_{t-1}^b(i)} \\ &= \text{market share of household } i \text{ in period } t - 1.\end{aligned}$$

We define the initial holdings of bonds for household i as

$$\begin{aligned}\bar{b}_t(i) &= (1 + i_{t-1}) \frac{w_{t-1}^b(i)}{\text{GDP}_t} \\ &= (1 + i_{t-1}) \frac{w_{t-1}^b(i)}{\sum_i w_{t-1}^b(i)} \frac{\sum_i w_{t-1}^b(i)}{\text{GDP}_{t-1}} \frac{\text{GDP}_{t-1}}{\text{GDP}_t}.\end{aligned}$$

As an interest factor for a positive (lending) position, we use an average of 6-year bond rates between the 4th year of period t and the 4th year of period $t - 1$. We add a spread for the borrowing rate. The spread is 2% for 1995 and 2.75% for 1968 and 1978, for reasons described in the calibration section IV.

Forecasting Income

The final step in our construction of the joint income and endowment distribution is to specify the marginal distribution of non-asset income. Here we make use of the fact that income is observed in period $t - 1$ in the SCF. We then assume that the transition between $t - 1$ and t is determined by a stochastic process for non-asset income. We employ the same process that agents in the model use to forecast their non-asset income, described in the next subsection. If the assumption were true, and if there were a large number of identical individuals in every cell, then our discretization implies that households in a cell should split up into nine different cells in the following period, with fractions provided by the probabilities of the income process. This is what we assume. As a result, the distribution of agents in period t is approximated by $9 \times 6 \times 9 = 486$ different cells. For each cell, we know the endowment of assets as well as income, and we have a set of population weights that sums to the total population.

Non-transiting households

The previous discussion has covered only households who transit from period t into period $t + 1$. We also need to take into account the creation and destruction of households between $t - 1$ and t . In years where successive SCFs are available, we calculate “birthrates” and “deathrates” for households directly by comparing these surveys. We assume that exiting households receive no labor income, but sell their assets and consume the proceeds, while entering households start with zero assets and the average labor income of their cohort. This is a simplified view that does not do justice to the many different reasons why households form and dissolve, and how wealth is passed along among households. However, we view it as a useful benchmark.

Time periods without two successive SCFs

For periods before 1980, the above strategy cannot be executed as is, because we do not have two consecutive SCFs. For the period 1965-70, the 1962 SCF is used to determine the initial endowment and income distribution. The only difficulty here is the adjustment of exiting and entering households. We use data from the Census Bureau on the evolution of household populations to gauge the size of exiters and entrants. The average labor income of the entering cohort is then estimated by multiplying per capita income of the young in the 1962 SCF by the growth rate of aggregate per capita labor income.

For the period 1975-80, we do not have SCF information for period $t - 1$. As for the 1960s, the updating of population weights is performed using Census data. To estimate the cross sectional distribution of endowments and income, we start from the 1962 distribution and its division of households into cells and modify cell holdings to obtain a new distribution. In particular, we calculate the unique distribution such that, for stocks, real estate, nominal assets, nominal debt and income, (i) aggregates match the aggregates from the FAUS for the period 1969-74, (ii) the ratio of holdings between individual members of any two cells is the same as in the 1962 SCF and (iii) the relative size of a cell within its age cohort is the same as in the 1962 SCF.

Condition (i) and (ii) imply that per capita holdings or income within a cohort changes in order

to account for differences in demographics while simultaneously matching aggregates. Conditions (ii) and (iii) imply that the cross section conditional on age is the same in the two years. The reason for using the 1962 distribution as the starting point rather than, say the 1989 distribution, is that the early 1970s aggregates – especially gross nominal assets – appear more similar to 1962 than to aggregates from the 1980s. Once we have a distribution of positions at the cell level for 1969-74, we proceed as above to generate an updated distribution for 1975-80.

C. Asset Supply

The endowment of the ROE sector consists of new equity issued during the trading period. The factor f^e states this endowment relative to total market capitalization in the model. We thus use net new corporate equity divided by total household holdings of corporate equity. We obtain the corresponding measure for housing by dividing residential investment by the value of residential real estate. The calibration of the model uses six-year aggregates. The initial nominal position of the ROE is taken to be minus the aggregate (updated) net nominal position of the household sector. Finally, the new net nominal position of the ROE sector in period t – in other words, the “supply of bonds” to the household sector – is taken to be minus the aggregate net nominal positions from the FAUS for period t .

B Baseline Expectations

This appendix describes agents’ expectations about returns and income in the future under the baseline scenario.

Non-Asset Income

We specify a stochastic process to describe consumer expectations about after-tax income. The functional form for this process is motivated by existing specifications for labor income that employ a deterministic trend to capture age-specific changes in income, as well as permanent and transitory components. In particular, following Zeldes (1989) and Gourinchas and Parker (2002), we assume that individual income Y^i is

$$Y_t^i = G_t A_t P_t^i U_t^i$$

which has a common component G_t , an age profile A_t , a permanent idiosyncratic component P_t^i and a transitory idiosyncratic component U_t^i .

The growth rate of the common component G_t is equal to the growth rate of aggregates, such as GDP and aggregate income, in the economy. It is common to specify the transitory idiosyncratic component as lognormally distributed

$$\ln U_t^i = N \left(-\frac{1}{2} \sigma_u^2, \sigma_u^2 \right),$$

so that U_t^i is i.i.d with mean one. The permanent component P_t^i follows a random walk with mean one. The permanent component solves

$$\ln P_t^i = \ln P_{t-1}^i + \varepsilon_t^i - \frac{1}{2} \sigma_\varepsilon^2,$$

where ε_t^i are normal shocks with zero mean and standard deviation $\sigma(\varepsilon_t^i)$. In our numerical procedures, we discretize the state process using Gauss–Hermite quadrature with three states.

We estimate the age profile G_t as average income in each age-cohort from the SCF:

$$\frac{1}{\#a} \sum_{i \in a} Y_t^i = A_t G_t \frac{1}{\#a} \sum_{i \in a} P_t^i,$$

with $\text{plim} \frac{1}{\#a} \sum_{i \in a} P_t^i = 1$. Table B.1 reports the profile relative to the income of the youngest cohort.

TABLE B.1: INCOME AGE PROFILE

29	35	41	47	53	59	65	71	77	88+
2.04	2.51	3.17	3.80	4.56	3.81	3.00	1.93	1.42	1.17

NOTE: Income age profiles estimated from the SCF. The numbers represent the average cohort income relative to the average income of the youngest cohort (≤ 23 years).

We obtain an estimate of the variance of permanent shocks by computing the cross-sectional variance of labor income for each cohort before retirement (<65 years) and then regressing it on a constant and cohort age. The intercept of this regression line is .78, while the annualized slope coefficient is .014.¹³ We thus set $\sigma_\varepsilon^2 = 1.4\%$ and time-aggregate this variance for our six-year periods. This number is in line with more sophisticated estimations of labor income processes, which tend to produce estimates between 1% and 2% per year.

Typical estimates of the variance of temporary shocks σ_u^2 are 2-10 times larger than those of the variance of permanent shocks. Moreover, several studies have shown that the variance of temporary shocks to log wages has increased since the 1970s. For example, Heathcote, Storesletten and Violante (2010) show that the variance of log wages increased from about .05 in the 1970s to .07 in 1995. To capture this increase in temporary income risk, we adopt their numbers, thus assuming that the variance of hours is constant. We set $\sigma_u^2 = .05$ for the years 1968 and 1978 and $\sigma_u^2 = .07$ for 1995.¹⁴ In all years, agents in the model assume that this variance is fixed forever. Finally, we determine the variance of the first draw of permanent income from the intercept of our regression line. For the earlier years, we scale down this initial variance by the relative change in the permanent component of income for the youngest agents in Heathcote et al. This is a simple way to accommodate changes in income due to education over time. Sensitivity checks have shown that the initial draw of permanent income does not matter much for the results, since it does not directly affect the portfolio problem.

¹³Of course, this simple approach uses only the cross-section and thus potentially confounds age and time or cohort effects. However, when we rerun the regression with SCF wages, using the similar sample selection criteria as Heathcote, Storesletten and Violante (2010), our results are close to what these authors find for 1995 from an analysis with panel data on wages.

¹⁴The fixed effects only matter for the updating of the income distribution, as explained in the previous section. The model's results are not sensitive to the magnitude of these effects. For example, the results based on our model are unchanged when we use the estimates provided by Heathcote, Storesletten and Violante (2010).

Most labor income studies focus on pre-retirement income. There are major challenges to obtaining variance estimates for retirement income. A key reason is that older households tend to experience large shocks to health expenditures, which are included as NIPA income if they are disbursed by health plans. These shocks contain both transitory and permanent components, such as long-term care. DeNardi, French and Jones 2010 provide estimates and show that these are quantitatively important for savings of the elderly (see also DeNardi, French and Jones, 2016 and Ameriks, Briggs, Caplin, Shapiro and Tonetti, 2020.) We therefore apply the above shocks to disposable income at any age, including retirement. Alternatively, we could introduce bequest motives, which would also prevent model-implied savings from dramatically dropping after retirement.

Returns and aggregate growth

We assume that consumers believe real asset returns and aggregate growth to be serially independent over successive six year periods. Moreover, when computing an equilibrium for a given period t , we assume that returns are identically distributed for periods beyond $t + 1$. We will refer to this set of beliefs – to be described below – as *baseline beliefs*. However, in our exercises we will allow beliefs for returns between t and $t + 1$ to differ. For example, we will explore what happens when expected inflation is higher over the next six year period. We discuss the latter aspect of beliefs below when we present our results. Here we focus on how we fix the baseline.

To pick numbers for baseline beliefs, we start from empirical moments. Table B.2 reports summary statistics on ex-post realized pre-tax real returns on fixed income securities, residential real estate and equity, as well as inflation and growth. These returns are measured over six year periods, but reported at annualized rates. Since we work with aggregate portfolio data from the FAUS, we construct returns on corporate equity and residential real estate directly from FAUS aggregates.

TABLE B.2: SUMMARY STATISTICS

r_t^b	r_t^h	r_t^e	π_t	g_t
Means				
2.68	4.81	8.51	4.01	2.19
Standard Deviations/Correlations				
3.24	-0.02	0.56	-0.04	0.33
-0.02	3.31	-0.04	-0.13	0.38
0.56	-0.04	22.87	-0.52	0.20
-0.04	-0.13	-0.52	5.60	-0.40
0.33	0.38	0.20	-0.40	1.31
Sharpe Ratios				
	0.45	0.27		

NOTE: The table reports annualized summary statistics of six-year log real returns. Below the means, the matrix has standard deviations on the diagonal and correlations on the off-diagonal. The last row contains the Sharpe ratios. The log inflation rate π_t is computed

using the CPI, while g_t is the log growth rate of GDP multiplied by the factor 2.2/3.3 to match the mean growth rate of consumption.

Baseline beliefs assume that the payoff on bonds $1/\pi_{t+1}$ is based on a (net) inflation rate $\pi_{t+1} - 1$ with a mean of 4% per year, and that the volatility of π_{t+1} is the same as the unconditional volatility of real bond returns, about 1.3% per year. To obtain capital gains from period t to $t + 1$, we take the value of total outstandings from the FAUS in $t + 1$, and subtract the value of net new issues (or, in the case of real estate, new construction.) To obtain dividends on equity in period t , we use aggregate net dividends. To obtain dividends on real estate, we take total residential housing sector output from the NIPA, and subtract materials used by the housing sector. For bond returns, we use a six year nominal interest rate derived by extrapolation from the term structure in CRSP, and subtract realized inflation, measured by the CPI. Here growth is real GDP growth.

The properties of the equity and bond returns are relatively standard. The return on bonds has a low mean of 2.7% and a low standard deviation of 3.2%. The return on stocks has a high mean of 8.5% and a standard deviation of 23%. What is less familiar is the aggregate return on residential real estate: it has a mean and standard deviation in between the other two assets. It is apparent that the Sharpe ratio of aggregate housing is much higher than that on stocks.

In principle, we could use the numbers from Table B.2 directly for our benchmark beliefs. However, this would not capture the tradeoff faced by the typical individual household. Indeed, the housing returns in Table B.2 are for the aggregate housing stock, while real estate is typically a non-diversified investment. It is implausible to assume that investors were able to pick a portfolio of real estate with return characteristics as in Table B.2 at any time over our sample period. Instead, the typical investor picks real estate by selecting a few properties local markets.

Existing evidence suggests that the volatility of house returns at the metro area, and even at the neighborhood or property level are significantly higher than returns at the national aggregate. For example, Caplin et al. (1997) argue that 1/4 of the overall volatility is aggregate, 1/4 is city-component, and 1/2 is idiosyncratic. Tables 1A and 1B in Flavin and Yamashita (2002) together with Appendix C in Piazzesi, Schneider and Tuzel (2007) confirm this decomposition of housing returns. As a simple way to capture this higher property-level volatility, we add idiosyncratic shocks to the variance of housing returns that have volatility equal to 3.5 times aggregate volatility. The resulting overall volatility of the return on an individual property is then 11%, consistent with the estimates in Landvoigt, Piazzesi and Schneider (2015). Finally, we assume that expected future real rents are constant. This ignores the volatility of real rent growth, which is small, at around 2% per year.

Taxes on investment

Investors care about after-tax real returns. In particular, taxes affect the relative attractiveness of equity and real estate. On the one hand, dividends on owner-occupied housing are directly consumed and hence not taxed, while dividends on stocks are subject to income tax. On the other hand, capital gains on housing are more easily sheltered from taxes than capital gains on stocks. This is because

many consumers simply live in their house for a long period of time and never realize the capital gains. Capital gains tax matters especially in inflationary times, because the nominal gain is taxed: the effective real after tax return on an asset subject to capital gains tax is therefore lower when inflation occurs.

To measure the effect of capital gains taxes, one would ideally like to explicitly distinguish realized and unrealized capital gains. However, this would involve introducing state variables to keep track of past individual asset purchase decisions. To keep the problem manageable, we adopt a simpler approach: we adjust our benchmark returns to capture the effects described above. For our baseline set of results, we assume proportional taxes, and we set both the capital gains tax rate and the income tax rate to 20%. We define after tax real stock returns by subtracting 20% from realized net real stock returns and then further subtracting 20% times the realized rate of inflation to capture the fact that nominal capital gains are taxed. In contrast, we assume that returns on real estate are not taxed.