Deutsche Bundesbank's 9th Spring Conference: "Microdata Analysis and Macroeconomic Implications"

Eltville, 27-28 April 2007

Martin Schneider
New York University and FRB Minneapolis

Monika Piazzesi
University of Chicago

„Asset Demand and the Distribution of Household Characteristics "

http://www.bundesbank.de

Asset Demand and the Distribution of Household Characteristics*

Monika Piazzesi
University of Chicago

Martin Schneider
NYU and FRB Minneapolis

April 2007

Abstract

This paper studies the effect of household heterogeneity on asset prices in an overlapping generations model with incomplete markets. The key new feature is that we focus on a particular trading period and use the Survey of Consumer Finances to directly measure the initial distribution of household characteristics. Conditional on age, changes in the distribution asset endowments and income have only a small effect on equilibrium prices, while changes in the distribution across age cohorts have significant effects. Moreover, idiosyncratic shocks to nontradable labor income and residential real estate are important determinants of household choices.

*Email addresses: piazzesi@uchicago.edu, martin.schneider@nyu.edu.

I Introduction

US households differ significantly in age, wealth and portfolio composition. They also experience idiosyncratic shocks, both to labor income and to asset positions, in particular to undiversified holdings of residential real estate. It is an open question whether modelling such heterogeneity explicitly is important for thinking about asset prices. Aggregation theorems that justify the standard representative agent approach to asset pricing hold only under restrictive assumptions. When agents have different preferences (for example because of age) and markets are incomplete (for example because agents cannot trade claims that are contingent on their income or the value of their property), those aggregation theorems do not apply. This paper follows a recent literature in asking whether heterogeneity is *quantitatively* important for asset prices. The main new feature of the analysis is that the distribution of household characteristics (in particular, age and asset positions) is measured from micro data and used directly as an input to the model.

We consider an OLG model with incomplete markets. Agents in the model trade the three major assets held by US households: equity, residential real estate, and fixed-income securities. The model accounts for aspects of the observed cross section of portfolio holdings. We use it to ask how aggregate asset demand and asset prices depend on two dimensions of heterogeneity: age and wealth. We find that conditional on age, changes in the distribution of asset endowments and income have only a small effect on equilibrium prices. In contrast, changes in the distribution across age cohorts have significant effects. We also consider the role of idiosyncratic shocks to labor income, as well as housing shocks. We find that both types of shocks are important for understanding asset prices.

We use the temporary equilibrium model of asset pricing proposed in Piazzesi and Schneider (2006). The model describes a particular trading period, which we take to be the mid 1990s. Successive waves of the Survey of Consumer Finances provide (*i*) the distribution of asset endowments and income at the beginning of the trading period, an exogenous input to the model, and (*ii*) the distribution of portfolios at the end of the trading period, an endogenous output used to evaluate the model. Individual asset demands are derived by solving household savings and portfolio choice problems, given initial endowments as well as expectations about future asset returns and income. Aggregate household sector asset demand is constructed by integrating individual demands under the initial distribution from (i).

Equilibrium prices are determined by equating household sector asset demand to the supply of assets, including new assets supplied to households by other sectors, measured from Flow-of-Funds data. The model derives equilibrium prices as a function of three exogenous inputs: the distribution of asset endowments and income, asset supply, and expectations. Our baseline exercise is calibrated so that the model provides a good match for equilibrium prices at the true distribution. We then explore the role of heterogeneity by evaluating the price function at counterfactual distributions that switch off some of the heterogeneity by age and wealth that is present in the data. We also explore the role of idiosyncratic shocks by eliminating them from the specification of households' expectations.

We present three main results. First, differences in asset endowments and income *within cohorts* have only small effects on asset prices. To be concrete, relative to our baseline economy that is based on the wealth distribution from the data, an economy with no within-cohort heterogeneity gives rise to a 8% higher stock market and a 3% higher value of the housing stock; the nominal interest rate increases by 20 basis points. The intuition for this result is similar to that behind the "approximate aggregation" result of Krusell and Smith (1998). Indeed, asset demand functions for most cohorts are almost linear in the relevant state variables – in our setting, both initial wealth (cash on hand) and the permanent component of non-asset income – at least in the support of the distribution of the state variables. Of course, in our setting the latter distribution is directly measured from the Survey of Consumer Finances, rather than being computed as part of the solution of the model.

Second, differences in asset endowments and income *across* cohorts have larger effect on asset prices. As one stark illustration, if all wealth were held by 47-53-year-old households, then the value of stocks and houses would be twice as large as in the baseline case. The reason is that middle aged cohorts have higher savings rates out of both asset wealth and non-asset income than households of other ages. This result says that measuring assets and income by age is much more important for our approach than measuring the distribution of wealth given age. It also suggests that demographic shifts or redistribution policies of income across generations may matter for asset prices.

Third, we show that the presence of idiosyncratic shocks is an important aspect of household beliefs in our model. Changing beliefs so that households ignore these shocks leads to lower asset values, as households try to reduce their precautionary savings. This effect is equally relevant for all types of assets. Moreover, reducing idiosyncratic shocks affects the portfolio composition, especially for young households. Indeed, it accentuates the role of human wealth as a riskless asset, and thereby increases the incentive for young households to invest their non-human wealth in risky ways. In our model, the presence of collateral constraints implies that risk-taking is best achieved by borrowing against real estate.

The idea behind the representative-agent approach to asset pricing is that the dynamics of prices is driven by the dynamics of macro data (and not, say, features of the wealth distribution), but that preferences can still be justified with reference to micro behavior. Researchers interested in asset prices therefore need not worry about the distribution of household characteristics. This paper suggests, in the context of our model, that this is only partially true. On the one hand, a model based on "cohort representative agents," whose preferences are consistent with individual behavior within the cohort, but who own all cohort wealth, would generate very similar prices.

On the other hand, preferences are sufficiently different across age groups that prices depend on the wealth distribution by age. It may be possible, of course, to reverse engineer a representative agent who gives rise to the same asset prices, if sufficient degrees of freedom with preferences are taken. However, our second result shows that the resulting preferences cannot be justified with reference to micro portfolio behavior in our economy. Instead, the backed out preference parameters would conflate preferences with the distribution of household characteristics. Researchers should thus take into account that distribution. In addition, our third result suggests that researchers should also take into account the distribution of idiosyncratic shocks.

The rest of the paper is organized as follows. Section II discusses related literature. Section III presents the model. Section IV 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 V discusses asset demand. It first discusses individual household asset demand and then studies aggregate demand with and without different layers of heterogeneity. Finally, Section VI reports the effects of heterogeneity on equilibrium prices.

II Related Literature

Strong aggregation results obtain in economies with complete markets and agents who have identical homothetic preferences. For example, in an endowment economy with these properties, asset prices are the same as in an economy with just one agent who has the same preferences, but owns the aggregate endowment. The reason is that agents in such an economy trade a rich set of assets to eliminate the effect of idiosyncratic shocks on consumption. In addition, homotheticity of preferences implies that the demand for assets is linear in wealth; aggregating over identical agents makes aggregate demand also linear in wealth. An observer studying such an economy can derive features of preferences (such as attitudes towards risk) from evidence on micro behavior (such as experiments or survey data), but can rely entirely on aggregate data to study the joint dynamics of asset prices and the rest of the economy. This observation provides the underpinnings for representative-agent asset pricing.

There is a large literature on models with incomplete markets and borrowing constraints. It is typically assumed that agents are infinitely-lived, have identical homothetic preferences and face idiosyncratic shocks to income. A limited asset structure can then destroy aggregation results, and in principle lead to large deviations from a representative agent benchmark (Constantinides and Duffie 1996). There are two effects. First, agents can no longer use a rich set of assets to trade away idiosyncratic risk. Second, the demand for assets need not be linear in wealth, even if preferences are identical and homothetic. This is because borrowing constraints introduce nonlinearities in demand. Nevertheless, there are a number of cases where strong aggregation results are still close to correct, because agents can use the limited number of asset to almost eliminate idiosyncratic risk and stay away from borrowing constraints. Loosely speaking, one will get closer to aggregation if agents are more patient, shocks to income have less permanent effects, aggregate risk is less important and borrowing is easier.[1]

Even if agents cannot trade away idiosyncratic shocks, it may still be true that the dynamics of asset prices does not depend on variables other than macro aggregates. This point was made by Krusell and Smith (1997, 1998) who develop an algorithm to compute stationary rational expecta-

[1] Analytical results on aggregation are in Levine and Zame (2002) and Krueger and Lustig (2007). Numerical studies of incomplete market economies are in Telmer (1993), Lucas (1994), Heaton and Lucas (1996), and Krusell and Smith (1997).

tions equilibria of heterogenous agents models. In their economies, the first moment of the wealth distribution is sufficient to capture the dynamics of equilibrium prices, a result they call "approximate aggregation." A difference to the strong aggregation results mentioned above is that moments of equilibrium asset prices may still depend on the distribution of idiosyncratic shocks. Intuitively, asset demand functions in the Krusell-Smith models – where agents have identical homothetic preferences – are approximately identical linear functions of wealth, at least in the support of the equilibrium wealth distribution. As a result, aggregate asset demand and market clearing prices depend only on aggregate wealth. However, the slope and intercept of an individual asset demand function can depend on how much idiosyncratic risk individuals face.

There are a few papers that numerically solve OLG models with aggregate risk and examines their implications for asset pricing. They typically focus on production economies. A key feature of OLG models is that one will commonly not assume identical homothetic preferences, but instead allow preferences to depend on age. The aggregation properties of OLG models are less well understood than those of infinitely-lived agent models. Storesletten, Telmer and Yaron (2004) have used Krusell-Smith algorithms to solve their model and found that the first moment of the wealth distribution is sufficient to capture aggregate dynamics. Krueger and Kubler (2004) propose an algorithm that treats the distribution of wealth across generations as a state variable and provide examples where this feature is important.

The present paper differs from previous studies in several ways. First, it considers rich asset and shock structures, motivated by our desire to calibrate to observed household asset positions. Existing studies focus on stocks and riskless bonds, and on shocks to labor income. In our setting, bonds are nominal, housing is present as an asset, and shocks to housing are a second source of idiosyncratic risk. Second, we do not compute stationary rational expectations equilibria (REE) of our model, but instead consider temporary equilibria for a single historical trading period. For this reason, our comparative statics exercises are not directly comparable to those in the literature. In our framework, it is natural to vary the distribution of endowments or that of idiosyncratic shocks while at the same time holding fixed expectations about future prices at values that are consistent with data. If equilibrium prices change in this experiment, we conclude that measuring this dimension of heterogeneity matters.

III 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 aggregate components of their housing dividends and labor income. They also face inflation risk when borrowing or lending because there is no riskfree asset. 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

$$C_t = c_t^{\delta} s_t^{1-\delta}. \tag{1}$$

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

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

where $U_{t+T} = C_{t+T}$. Here σ determines the intertemporal elasticity of substitution for deterministic consumption paths, 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 θ_t^h units of real estate, he obtains a dividend $(p_t^s - m)\theta_t^h =: d_t^h \theta_t^h$. Consumers form expectations about future real returns on housing and rental prices $\left\{R_\tau^h, p_\tau^s\right\}_{\tau>t}$. In particular, the real return on housing is affected by aggregate shocks $\overline{R}_t^h$ and idiosyncratic shocks ε_t, so that we can write

$$R_t^h = \overline{R}_t^h \varepsilon_t. \tag{3}$$

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 consumers has been a net borrower in the past. In period t, consumers can buy or sell bonds at a price q_t. A consumer expects every bond bought to pay $1/\pi_{t+1}$ units of 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 $(1+\xi)/\pi_{t+1}$ units of numeraire in period $t+1$, where $\xi > 0$ 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 ex-dividend value of all real estate owned by the consumer. For periods $\tau > t$, consumers form expectations about the (random) real return on bonds $\left\{R^b_\tau\right\}$. They believe that $R^b_\tau = 1/q_{\tau-1}\pi_\tau$ is the (ex post) real lending rate, and that $R^b_\tau(1+\xi)$ 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. Consumers believe their income evolves through time according to the process

$$y_\tau = \underbrace{G_\tau A_\tau P_\tau}_{\hat{y}_\tau} u_\tau \tag{4}$$

which has a permanent component $\hat{y}_\tau$ and a transitory idiosyncratic component u_τ with mean $E[u_\tau] = 1$. The permanent component $\hat{y}_t$ consists of a common component G_τ, an age profile A_τ, and a permanent idiosyncratic component P_τ with mean $E[P_\tau] = 1$. (This income specification follows Zeldes 1989 and Gourinchas and Parker 2001.)

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

[3] One 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 $(1+\xi)/\pi_{t+1}$ units of numeraire from every borrower (bond seller), but pay only $1/\pi_{t+1}$ to every lender (bond buyer), keeping ξ/π_{t+1} for itself. We do not model the financial intermediary explicitly since we only clear markets in period t.

Budget Set

The consumer enters period t with an endowment of houses and equity $\left(\bar{\theta}_t^h, \bar{\theta}_t^e\right)$ 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

$$\bar{w}_t = (p_t^h + d_t^h)\bar{\theta}_t^h + (p_t^e + d_t^e)\,\bar{\theta}_t^e + \bar{b}_t + y_t. \tag{5}$$

To allocate this initial wealth to consumption and purchases of assets, the consumer chooses a plan $a_t = \left\{c_t, s_t, \theta_t^h, \theta_t^e, b_t^+, b_t^-\right\}$, 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, $b_t^+ \geq 0$ 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

$$c_t + p_t^s s_t + w_t = \bar{w}_t, \tag{6}$$

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

$$\begin{aligned} & c_\tau + p_\tau^s s_\tau + w_\tau^h + w_\tau^e + w_\tau^{b+} - w_\tau^{b-} \\ & = R_\tau^h w_{\tau-1}^h + R_\tau^e w_{\tau-1}^e + R_\tau^b w_{\tau-1}^{b+} - R_\tau^b\,(1+\xi)\,w_{\tau-1}^{b-} + y_\tau. \end{aligned} \tag{7}$$

We denote the consumer's overall plan by $a = \left(a_t, \{a_\tau\}_{\tau=t+1}^{t+T}\right)$. This plan is selected to maximize utility (2) subject to the budget constraints (5)-(7) and the collateral constraints.

From the budget constraint (7), the solution to the optimization problem depends on the consumer's expectations about real returns $\left(R^h_\tau, R^e_\tau, R^b_\tau\right)$ and real income y_τ for periods $\tau > t$. These expectations are based on processes for returns and income that are affected by two types of shocks: aggregate shocks and uninsurable idiosyncratic shocks. In particular, housing returns $R^h_\tau = \overline{R}^h_\tau \varepsilon_\tau$ and income $y_\tau = G_\tau A_\tau P_\tau u_\tau$ are affected by idiosyncratic shocks; the shocks ε_τ and u_τ are i.i.d. in logs, while $\ln P_t$ is a random walk driven by an i.i.d. shock. The remaining processes $\left(\overline{R}^h_\tau, R^e_\tau, R^b_\tau, \Delta \ln G_\tau\right)$ are only affected by aggregate shocks.

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 (5). However, they do not make any savings or portfolio decisions. Instead, they simply purchase numeraire and housing services in the period t goods markets to maximize (1) subject to the budget constraint

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

B. Equilibrium

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

The Rest of the Economy

To close the model and regulate the supply of assets exogenous to the household sector, we introduce a rest-of-the-economy (ROE) sector.[4] 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^e_t trees and f^h_t houses in period t. Here f^e_t could be negative to represent repurchases of shares by the corporate sector. In addition, the ROE enters period t with an outstanding debt of $\bar{B}_t$ units

[4]The ROE sector allows us to accommodate household sector savings–deviations of consumption from household sector income–that are usually ignored in endowment economy models. More details are in Piazzesi and Schneider (2006).

of numeraire, and it raises D_t units of numeraire by borrowing in period t. The surplus from these activities is

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

If C_t^{ROE} is positive, it is consumed by the ROE sector. More generally, the ROE is assumed to have "deep pockets" out of which it pays for any deficit if $C_t^{ROE} < 0$.

Aggregate Asset Supply

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

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

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:

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

Equilibrium

An equilibrium consists of a vector of prices for period t, $\left(p_t^h, p_t^e, q_t, p_t^s\right)$, 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 $a\,(i) = \left(a_t\,(i)\,, \{a_\tau\,(i)\}_{\tau=t+1}^{t+T(i)}\right)$ given consumer $i's$ endowment, planning horizon, and expectations about future prices and returns;

(2) markets for all assets and goods clear:

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

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.

The equilibrium concept we use here is Grandmont's (1977, 1982) *temporary* equilibrium which does not place restrictions on prices beyond date t itself. Our empirical implementation follow the literature on portfolio choice, where exogenous processes for returns and income are standard and estimated from historical data. However, we go beyond portfolio choice models in that we explore how equilibrium *prices* change when investors' characteristics vary in a controlled way.

IV 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.

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.

We focus on the period 1992-97 and construct asset and income endowment distributions with data on asset holdings from the precursor period 1986-91. Micro data is not available at high frequencies. To capture the wealth and income distribution during the period, we have chosen these intervals so that the 4th year in each of these periods coincides with a Survey of Consumer Finances (SCF); we thus use surveys from 1989 and 1995.

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 Flow of Funds Accounts (FFA) 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 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) pensions 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 FFA/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_t^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 f_t^h consists of residential investment, divided by the value of residential real estate.

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.

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 FFA for period t.

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 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

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

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 the stochastic process (4) for non-asset income, which is the same process that agents in the model use to forecast their non-asset income. This approach allows to capture the correlation between income and *initial* asset holdings that is implied by the joint distribution of income and wealth.

C. Expectations, Preference and Credit Market Parameters

Non-Asset Income

We estimate the deterministic income age profile as average income in each age-cohort from the SCF. Table 1 reports the profile relative to the income of the youngest cohort. 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.[5] We thus set $\text{var}(\ln P_t) = 1.4\%$ 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 $\text{var}(\ln u_t)$ are 2-10 times larger than those of the variance of permanent shocks $\text{var}(\Delta \ln P_t)$. We adopt the estimates by Heathcote et al. (2004) who show that the variance of temporary shocks to log wages is $\text{var}(\ln u_t) = .07$ in 1995.[6] Finally, we determine the variance of the first draw of permanent income from the intercept of our regression line.

[5] Of 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 Storesletten et al. (2004), our results are close to what these authors find for 1995 from an analysis with panel data on wages.

[6] The 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 et al. (2004).

Table 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).

Most labor income studies focus on pre-retirement income. There are major challenges to obtaining variance estimates for retirement income. For example, 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 (see the estimates reported in Appendix A, Skinner, Hubbard and Zeldes 1994). Since these shocks are hard to measure at the household level, we could try to ignore their variances and assume that household receive a safe stream of income during retirement. However, this implies that precautionary savings drop dramatically as soon as the household enters retirement in the absence of such shocks. This prediction is not consistent with the household savings data from the SCF. For this reason, we apply the above shocks to income at any age, including 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 $t+1$. 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 house return faced by an individual household.

Table 2 reports annualized summary statistics. Since we work with aggregate portfolio data from the FFA, we construct returns on corporate equity and residential real estate directly from FFA aggregates. 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 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 FFA 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.

Table 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.

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%. The aggregate component of the return on residential real estate has a mean and standard deviation in between the other two assets. 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 ε_t to the variance of housing returns that have volatility equal to 3.5 times aggregate volatility.[7] 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.

Preference parameters

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 with historical Sharpe ratios on stocks and housing. To avoid this counterfactual implication, we assume that agents view long-lived assets as riskier than indicated by their historical moments. This idea of "low aversion against high perceived risk" can be captured by scaling the historical return variances from Table 2 with a fixed number. This scaling can be interpreted as a consequence of Bayesian learning about the premium on equity and housing. We select a factor of 3, which leads to match the aggregate portfolio weights for 1995 reported in Table 3 almost exactly.

Table 3: 1995 Data and Baseline Beliefs

experiment	wealth/ GDP	portfolio weights bonds	housing	stocks	lend./ GDP	borr/ GDP	PD ratios housing	stocks	interest rate nominal
(1) —1995 data —	2.51	.15	.59	.26	.70	.31	20.4	23.9	.061
(2) baseline	2.51	.15	.60	.25	.70	.31	20.8	22.7	.061

Note: The first row reports the aggregate portfolio weights on bonds, housing and stocks from Figure **??**; the gross borrowing and lending numbers from Section **??**, the wealth-to-GDP ratio from Figure **??**; the price-dividend ratios for housing and stocks together with the nominal 6-year interest rate. The second row report the results computed from the model for 1995 with baseline beliefs.

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 2, but are characterized by high risk aversion, $\gamma = 25$, and high discounting, $\beta = \exp(-0.07 \times 6)$.

[7] Since the volatility of housing is measured imprecisely, we chose the precise number for the multiplicative factor such that the aggregate share of housing in the model roughly matches the FFA data. The resulting factor is 3.5, close to the rule-of-thumb factor of 4.

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]

Credit Market Parameters and Taxes

It remains to select parameters to capture taxes on investment as well as consumers' opportunities for borrowing. We assume a 2% per year spread between borrowing and lending interest rates. In addition, we select the borrowing constraint parameter $\phi = .8$. This implies a maximal loan-to-value ratio of 80%, where "value" is the ex-dividend value of the house. We assume proportional taxes on stock returns, 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.

V Household Behavior

In this section, we consider savings and portfolio choice in the cross section of households. The initial distribution of asset endowments and income is derived from the 1989 SCF, as discussed in Section IV. We also compare the predictions of the model for the cross section of households in 1995 to actual observations from the 1995 SCF.

A. Individual Household Demand

The household problem is to maximize utility subject to the budget constraints (6)-(7). Since we are assuming i.i.d. returns and a constant rental price, there are only two state variables: household wealth at the beginning of a trading period, $\bar{w}_t$ (which consists of asset wealth plus non-asset income, as in equation (6)) and the permanent component of non-asset-income $\hat{y}_t$ in equation (4). The latter is needed to forecast future non-asset-income. Since utility is homothetic and all constraints are

[8] Yet 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).

linear, a standard argument implies that the problem can be rewritten in terms of a single state variable, the ratio of initial wealth to the permanent component of income $\bar{w}_t/\hat{y}_t$. For simplicity, we refer to the state variable as the wealth-to-income ratio.

Figure 1 plots household decision rules as a function of the wealth-to-income ratio for different age groups. The wealth-to-income ratio is measured along the horizontal axis. Since initial wealth contains income, households with no assets will have wealth-to-income ratios around one. Asset demand is measured along the vertical axis, and is normalized by permanent income. The top left panel shows total savings, or terminal wealth, while the other three panels show houses, equity and bond positions, all as fractions of permanent income. In each panel, black is for 71-year old households, dark gray is for 47-year olds, and light gray is for 35-year olds. These age numbers refer to the upper range of the age group. For example, "71-year old households" refers to households that are 66-71 years old.

Savings

To get intuition about the optimal decision rules, it is helpful to distinguish the savings and the portfolio decision. The goal of the savings decision is to smooth consumption over time. Consider first an agent with a low wealth-to-income ratio. This agent has little cash on hand at date t (low $\bar{w}$), relative to the non-asset income that he expects to obtain in the future (because permanent income $\hat{y}$ is high today). In order to smooth consumption, this agent would like to borrow against future labor income, but he cannot do so because the collateral constraint precludes negative net worth. Instead, an agent with low $\bar{w}/\hat{y}$ will save nothing and spend all cash on hand on consumption. This explains the flat portion of the optimal savings function.

Now consider an agent with very high $\bar{w}/\hat{y}$. If initial wealth is very high, consumption over the rest of the lifecycle will be financed almost exclusively out of initial wealth, and not out of future non-asset income. The agent will thus behave similarly to a "rentier" agent who has no non-asset income. In a world of i.i.d. returns, a rentier simply consumes some age-dependent fraction of wealth every period. A rentier's savings function (conditional on age) is therefore linear in $\bar{w}$. This is why at high values of $\bar{w}/\hat{y}$, where agents are very similar to rentiers, their savings function is nearly linear. Finally, for intermediate values of $\bar{w}/\hat{y}$, consumption smoothing suggests savings a smaller fraction out of every available dollar at date t, the lower is $\bar{w}/\hat{y}$. With an increasing savings *rate* the savings

Figure 1: Asset holdings and terminal wealth, both as fractions of permanent income, plotted against the initial wealth-to-income ratio. In each panel, black stands for 71-year old households, dark gray is for 47-year olds, and the light gray is for 35-year olds. Age groups are identified by maximum age in the cohort.

function becomes increasing and convex in $\bar{w}/\hat{y}$.

Comparison of the different lines in Figure 1 illustrates how savings depends on age. There are two relevant effects. First, 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. This effect matters most when labor income is not very important, that is, when the wealth-to-income ratio is high. For a wealth-to-income ratio of 5, the 35-year-old households save more than the 71-year-old group. For even higher $\bar{w}/\hat{y}$ (not shown), the young "rentier" households will eventually save also more than the 47-year-old group.

A second effect is that the age profile of non-asset income is hump-shaped, so that middle-aged investors can rely more on labor income for consumption than either young or old investors. This effect makes middle-aged investors save relatively more than other investors, and very young investors save relatively little. It matters most for investors with low initial wealth, who rely more on non-asset income for future consumption. For example, at values of $\bar{w}/\hat{y}$ between 1 and 2, the savings rate of the middle-aged is highest, and the savings rate of the old is higher than that of the young.

Portfolio Choice

The goal of the portfolio decision is to optimize the risk-return tradeoff for total wealth, which comprises both human and nonhuman wealth. Here the key principle is that human wealth is treated much like a riskless asset. This is due to the low correlation of implicit returns on human wealth with other asset returns – especially when there are idiosyncratic shocks – and to the favorable Sharpe ratio of human wealth returns. The main implication is that households with little nonhuman wealth relative to human wealth – in our setting, households with low wealth-to-income ratios – will choose a more risky portfolio for their nonhuman wealth. Versions of this basic effect are common in the literature (see, for example Jagannathan and Kocherlakota 1996, Heaton and Lucas 2000).

In our setting, an important way to build risky portfolios of nonhuman wealth is to borrow against real estate. As a result, the desire of low-wealth-to-income households to take risks translates into a desire to take out mortgages and buy housing. While holding equity is also a way to increase risk, equity is not as desirable for low wealth-to-income households because it cannot serve as collateral. In addition to the wealth-to-income ratio, age also reduces the willingness to take on risk. The hump-shaped profile for non-asset income implies that the implicit return on human capital becomes worse with age. This makes human capital less attractive and weakens the desire to take risk with nonhuman wealth.

To understand the asset demand functions in Figure 1, it is helpful to focus first on the lower left panel, which depicts bond positions. Beyond the far left region of the state space where all asset positions are zero, there are three further regions. First, as $\bar{w}/\hat{y}$ increases to a point where savings are positive, bond positions become negative. As $\bar{w}/\hat{y}$ increases further, there is a second region in which households stay out of the bond market. This nonparticipation region is caused by the credit spread: here the borrowing rate is perceived as too high to borrow, but the lending rate is not attractive enough to lend. Finally, for high values of $\bar{w}/\hat{y}$, households hold positive amounts of bonds.

The first region represents low wealth-to-income households leveraging and buying houses. The effect is seen most clearly for the youngest agents: at values of $\bar{w}/\hat{y}$ around 1, these agents do not invest in equity, but have substantial housing positions that are financed by borrowing. For older agents, the effect is weaker since human capital is not as attractive. In the third region, at high wealth-to-income ratios, human capital is sufficiently small relative to nonhuman wealth that it makes sense to invest in bonds. Given $\bar{w}/\hat{y}$, the desire to own bonds is higher for older and middle-aged agents. For all age groups and assets, demand functions eventually become linear. This is because, in a world with i.i.d. returns, rentier agents – and agents sufficiently similar to rentiers – optimally invest fixed fractions of savings in each of the three assets.

B. Cohort demand and heterogeneity within cohorts

We now ask whether cohort asset demand depends on the distribution of wealth and income within a cohort. We compare cohort asset demands for two initial distributions of endowments and income. The first is the distribution from the data, which is an input to our baseline exercise. The second distribution assigns all asset endowments and income for agents of a given cohort to a single household. We emphasize that the only difference between the two asset demands is in the initial distribution of endowments and income. In particular, return expectations are held at their baseline values and households face the same amount of idiosyncratic risk in both exercises. Moreover, market prices are held fixed at their baseline values; we consider the effect of the initial distribution on equilibrium prices in Section V below.

The results are summarized in Figure 2. There is again a top left panel that shows total savings as well as panels for the three individual assets. Every panel contains three lines. The dotted dark line is the cohort aggregate position as measured from the SCF data, stated as a percent of annual GDP. The solid dark line is the model-implied demand for the baseline calibration. Finally, the solid light line represents the model-implied demand when every cohort consists of a single household who starts out with the aggregate cohort endowments and income.

Positions in both model experiments exhibit the same general shape as the data. First, cohort savings are nonnegative and hump-shaped in age. The model generates nonnegative savings because of the collateral constraint. It generates a hump-shape in savings because both cohort income and

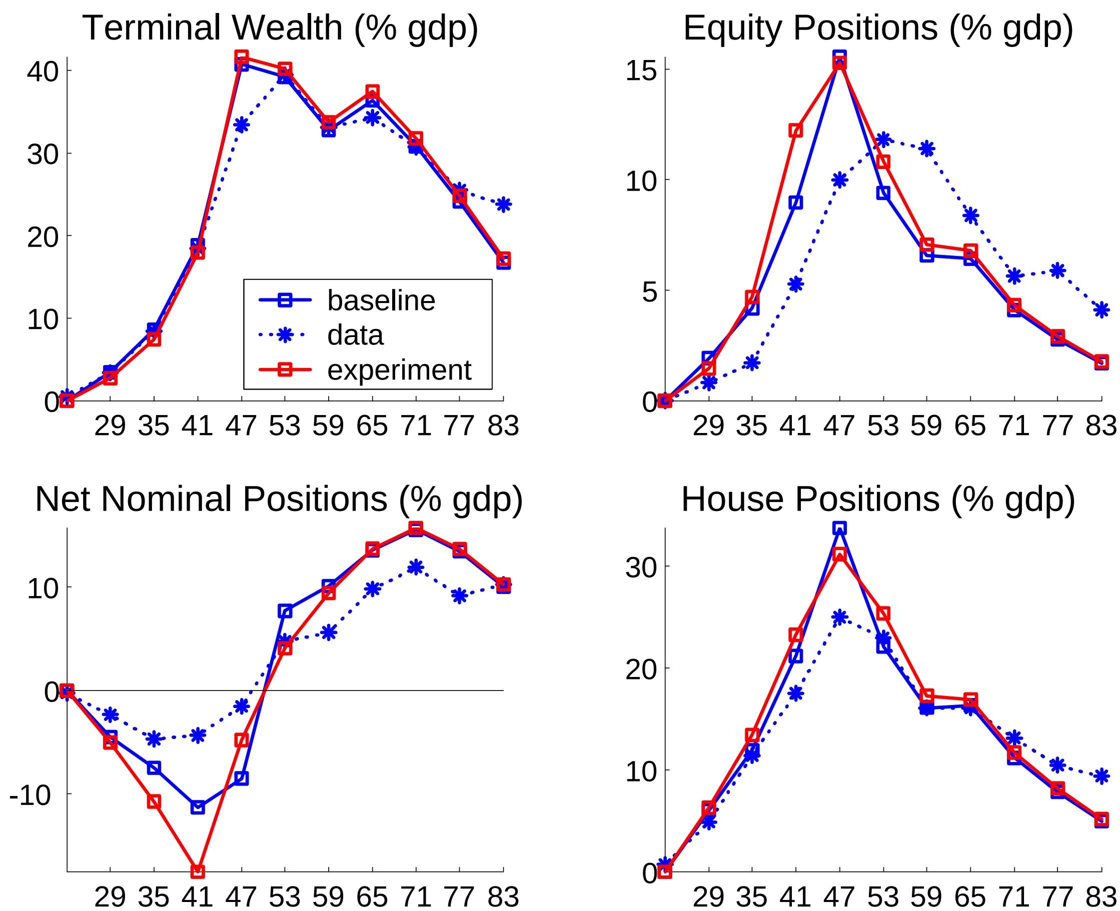

Figure 2: Asset positions by cohort in the data, implied by the baseline model, and implied by the model with only one household per cohort ("experiment"). The positions are in percent of GDP.

the savings rate out of initial wealth are hump-shaped in age. Second, young households tend to be net borrowers in the bond market, while older households are net lenders. The model generates this because younger households, who typically have lower wealth-to-income ratios, build more risky asset positions via leverage. Finally, house and equity positions, which are not allowed to be negative, reflect the overall hump shape of total savings.

After we switch off heterogeneity in wealth and income, total savings are lower for younger households (aged under 41), and they are higher for older households. Portfolio behavior changes generally in the direction of more risk-taking: for younger households this is due to more leverage and for middle-aged households it is due to less bond-holdings relative to holdings of the more risky stocks and houses. The biggest changes appear in the portfolio behavior of 29-41 year old households. For these cohorts, borrowing in the credit market almost doubles, and the borrowed funds are invested in equity and housing.

Approximate aggregation for savings

To obtain intuition for these results, consider savings behavior within a cohort. All agents in a cohort have identical homothetic preferences. If markets were complete, or, more generally, if agents' endowments were in the span of tradable assets, portfolio demand would be linear in initial wealth. Aggregate demand would then depend only on aggregate initial wealth. As we have seen in Figure 1 above, individual asset demands from our model are not linear over the entire state space. However, cohort demand will not depend strongly on the wealth distribution as long as savings functions are approximately linear *in the support* of the initial distribution of the wealth-to-income ratio.

To be concrete, fix one cohort and let the index i refer to households in the SCF, or cells in our approximate distribution of households. Let $g(i)$ denote the weighting function for households that is used to compute cohort aggregates. For example, aggregate cohort initial wealth is $\bar{W} = \sum_i \bar{w}^i g(i)$ and so on. Since preferences are identical within the cohort, the total savings of household i can be written as $s\left(\bar{w}^i/\hat{y}^i\right)\bar{w}_i$, where s is a savings rate function that is specific to the cohort. Suppose that the function s is such that

$$s\left(\bar{w}^i/\hat{y}^i\right)\bar{w}^i/\hat{y}^i \approx \alpha + \beta\left(\bar{w}^i/\hat{y}^i\right).$$

in the support of g. The top left panel of Figure 1 suggests that this shape, with $\alpha < 0$ and $\beta > 0$, could be a good approximation at least for old households. Roughly, rich households with a higher permanent component of income save less overall, but they save the same fraction of every additional unit of wealth.

Cohort savings now becomes

$$\begin{aligned} W &= \textstyle\sum_i s\left(\bar{w}^i/\hat{y}^i\right)\bar{w}^i\; g(i) \\ &\approx \textstyle\sum_i \left(\alpha + \beta\left(\bar{w}^i/\hat{y}^i\right)\right)\; \hat{y}^i\, g(i) \\ &= \textstyle\sum_i \left(\alpha\,\hat{y}^i + \beta\,\bar{w}^i\right)\; g(i) \\ &= \alpha\, Y + \beta\,\bar{W}, \end{aligned} \tag{8}$$

where Y is aggregate cohort income (which equals aggregate cohort permanent income by (4), $Eu_t^i = 1$, and the fact that the law of large number applies within the cohort). Cohort asset demand can thus be understood by considering jointly the individual demand functions of Figure 1 and the distribution of the wealth-to-income ratio in the cohort population.

The top panel of Figure 3 reproduces the top left panel of Figure 1, which shows total savings relative to permanent income – from now on, the savings-to-income ratio – as a function of the wealth-to-income ratio. The bottom panel reports, for the same selected cohorts, the cumulative distribution functions for the wealth-to-income ratio, after an adjustment that multiplies the weight on every household i by a factor $\hat{y}^i/Y$. The adjustment implies that integrating the savings function for a given cohort in the top panel under the density implied by the cdf for that cohort in the bottom panel delivers the ratio W/Y of aggregate cohort savings to aggregate cohort income.

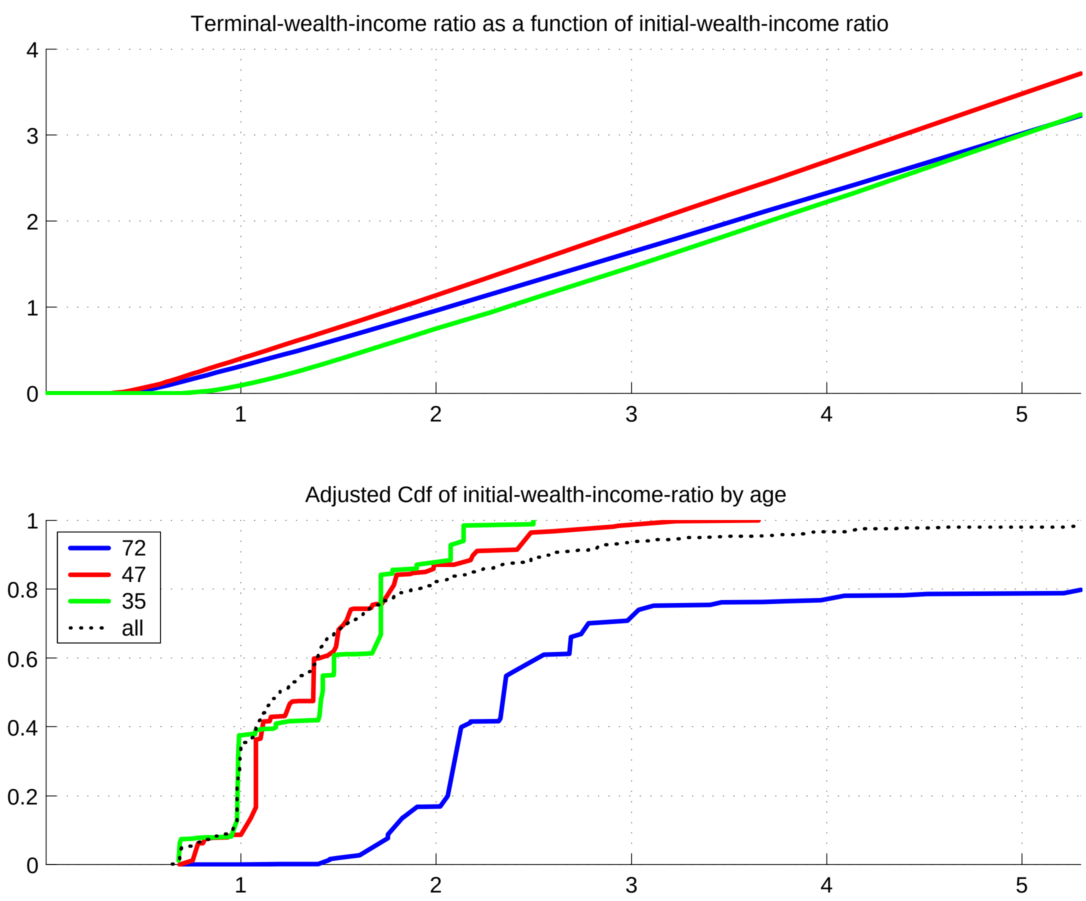

Figure 3: The upper panel shows savings-to-income as a function of wealth-to-income. The lower panel shows the adjusted cdf of households over these wealth-to-income levels.

The figure shows that, for all cohorts, many households are located in the region of the state space where the savings function is linear. As a result, the deviations in aggregate savings between the two experiments in Figure 2 are nowhere larger than a few percentage points of GDP. Conditional on age, our model thus gives rise to what Krusell and Smith (1998) have called an "approximate aggregation" result.

Nonlinear savings functions can cause demand without income and wealth heterogeneity to be either higher or lower than demand in the baseline model. One effect is that the savings function is convex in the wealth-to-income ratio. This implies that the cohort savings-to-income ratio without wealth and income heterogeneity is always below the mean savings-to-income ratio in the population. If income is the same across agents, so that differences in the savings-to-income ratio are only driven by savings itself, then it follows that cohort savings without heterogeneity is lower than cohort savings in the baseline case. The result continues to be true when the dispersion of permanent income is not too high. This explains why young cohorts save less when wealth and income are aggregated within the cohort.

A second effect becomes relevant if agents with higher savings have relatively low income. This is relevant for older cohorts, where some "worker" households have high income and little savings, while other "rentier" agents have a lot of savings and little income. In this case, aggregating income and wealth within the cohort generates an aggregate household with a high wealth-to-income ratio who actually saves more than the cohort saves in the baseline model. In Figure 2, this effect is most pronounced for middle aged households nearing retirement.

Lack of aggregation for bond demand of the young

Figures 4 and 5 juxtapose the distribution of the wealth-to-income ratio and demand functions for the individual assets for the relatively young 35-year-old cohort and the relatively old 71-year-old cohorts, respectively. In both figures, thick solid lines represent the asset demand functions already shown in Figure 1. The thin solid lines on top of the shaded areas are the adjusted cdfs of the wealth-to-income-ratio from Figure 3. As for aggregate savings, differences between demand with and without wealth and income heterogeneity requires nonlinearity in the support of the distribution.

The most obvious example of nonlinearity is in the bond demand of the young in Figure 4. It is

Figure 4: Young households' asset demand (35-year olds). The thin lines represent adjusted cdfs and have arbitrary scales in this graph for better visual comparability.

apparent from the figure that giving all wealth and income to an agent in the middle of the support will tend to increase borrowing of the cohort. This is precisely what happens for the young cohorts in Figure 2. Nonlinearities in housing and stock demand are less pronounced; for this reason the differences between the two exercises are smaller for those assets. For the old generations, there are no apparent nonlinearities, which accounts for the fact that approximate aggregation obtains for these assets in Figure 2.

C. Differences in demand across cohorts

So far, we have focused on aggregation within a single cohort. However, the argument using equation (8) is also relevant for thinking about aggregate demand by the whole household sector. In particular, if all cohort demands can be represented, or well approximated, by linear functions with the same coefficients α and β, then changes in the distribution of wealth and income *across* cohorts do not

Figure 5: Old households' asset demand (71-year olds). The thin lines represent adjusted cdfs and have arbitrary scales in this graph for better visual comparability.

affect aggregate asset demand. In other words, the distribution across cohorts would not matter at all if the savings functions in Figure 1 were linear functions with the same intercept and the same slope, at least in the support of the distribution of households.

We have already seen that, while savings functions are not linear in all relevant regions of the state space, the approximation by a linear function is not bad. This suggests the following way to visually compare savings functions across cohorts. We calculate, for every cohort, the tangent to the graph of the savings function $s(x)x$ at the point where the wealth-to-income ratio is equal to aggregate wealth divided by aggregate income. This delivers a set of coefficients (α, β) for every cohort. The implied linear savings functions $\max\left\{\alpha + \beta \frac{\bar{w}}{y}, 0\right\}$ provides a reasonable approximation to the actual cohort savings function.[9] Figure 6 plots β as well as the sum $\alpha + \beta$.

[9] More specifically, aggregate cohort savings under the constructed savings function is within 2% of GDP to true cohort savings in the baseline case.

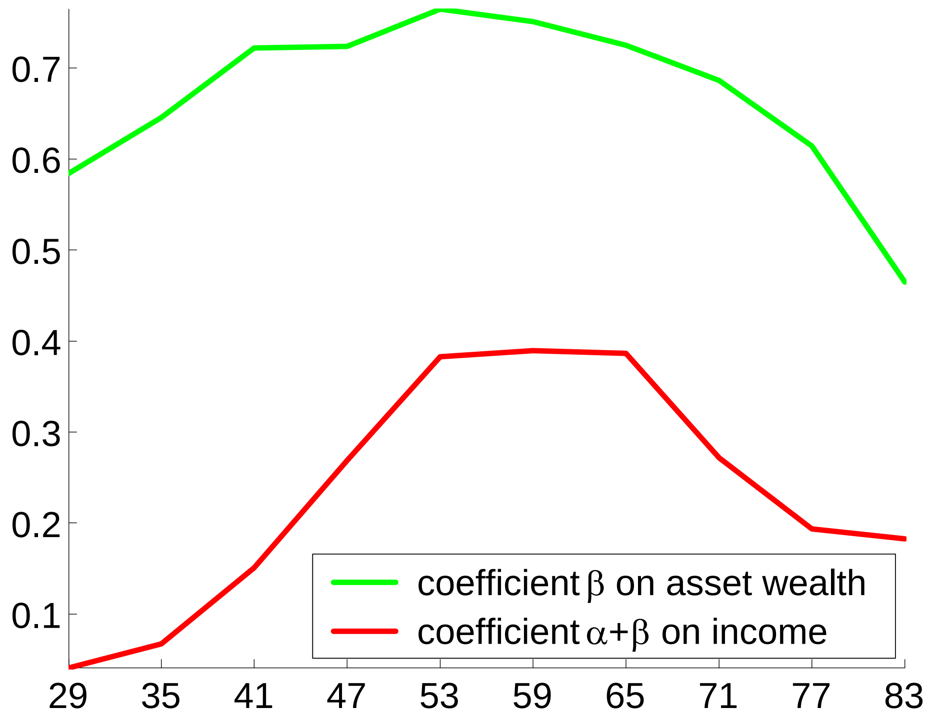

Figure 6: Coefficients on income and asset wealth when cohort savings are approximated by linear functions.

The sum $\alpha + \beta$ is of represents the marginal propensity of the cohort to save out of cohort non-asset income, whereas β is the propensity to save out of asset wealth. Indeed, (8) together with the budget constraint (5) imply that

$$\begin{aligned} W_t &= \alpha Y_t + \beta \bar{W}_t \\ &= (\alpha + \beta) Y_t + \beta \left(p_t^h + d_t^h + p_t^e + d_t^e + \bar{B}_t \right) \\ &= (\alpha + \beta) Y_t + \beta \bar{W}_t^a, \end{aligned}$$

where $\bar{W}^a$ is the value of assets owned by the cohort at the beginning of the period. The propensity to save out of non-asset income is lower ($\alpha < 0$), because higher non-asset income at date t signals higher non-asset income in the future, which reduces the need for savings.

We can read off Figure 6 how changes in the distribution of either non-asset income or asset wealth across cohorts can change aggregate savings. Both propensities to save differ across cohorts. Changes in the distribution of income have larger effects, not only because the propensity to save out of non-asset income differs more across cohorts, but also because income is a more important component of initial wealth than asset wealth. Indeed, aggregate income is about .6 GDP, whereas aggregate initial asset wealth is about .4 GDP (here we are working with 6 year periods).

To get an idea about magnitudes, consider the following rough example calculations. If all income were earned by households between 53 and 65, then aggregate savings out of income would be about $(.4)(.6) = .24$ GDPs, almost 4 times larger than if all income were earned by households under 41, where we would get only $(.1)(.6) = .06$ GDP. Changes in the assignment of asset endowments have smaller effects. For example, if all asset wealth were held by agents between 41 and 65, then aggregate savings out of asset wealth would be $(.7)(.4) = .28$ GDPs, whereas if all asset wealth were held by by agents younger than 29 or older than 77, then it would be at most $(.6)(.4) = .24$. Both changes thus have significant effects on aggregate savings. We explore the impact of these changes further when we consider asset pricing in the next section.

D. The role of idiosyncratic shocks

In our model, the fact that the household sector consists of a large number of different households is reflected not only through a non-degenerate wealth and age distribution, but also in the types of risk households perceive for the future. In particular, there are idiosyncratic shocks to non-asset income shocks and to the return on a household's housing position. In this subsection, we ask what happens to cohort demand when these idiosyncratic risks are turned off.

We report the result in two figures that have the same structure as Figure 2. Figure 7 shuts down idiosyncratic shocks to non-asset income. There are four panels for savings and the three individual assets, and we plot demand relative to annual GDP for the data, the baseline case as well as an experiment where households face no idiosyncratic shocks to labor income. Except for elderly households, savings drop substantially. This is due especially to a reduction of equity holdings. Housing positions are lower for the youngest households, but middle-aged households retain a strong demand for housing, accompanied by substantially higher mortgage borrowing than in the baseline case.

The intuition for this result can be obtained from Figure 1 and our discussion of individual

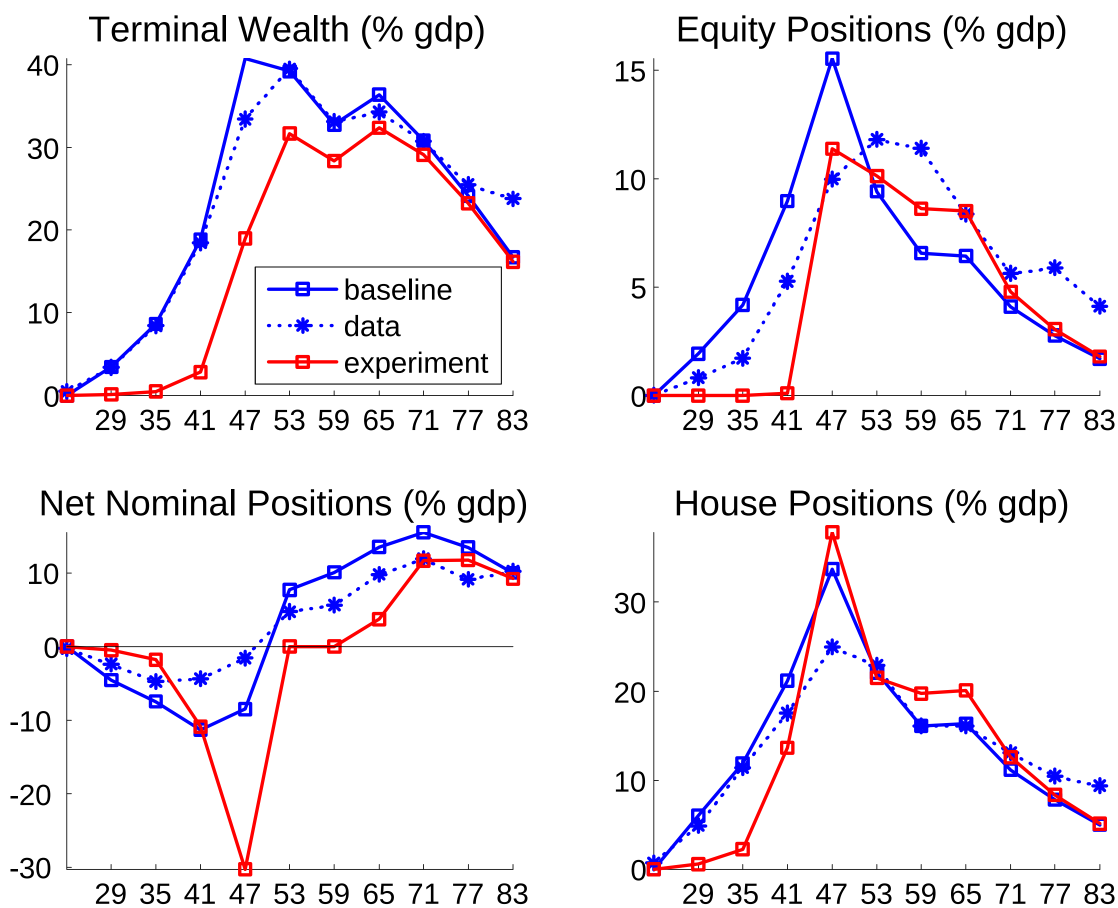

Figure 7: Asset positions by cohort in the data, implied by the baseline model, and implied by the model without idiosyncratic labor income risk ("experiment"). The positions are in percent of GDP.

demand. The removal of idiosyncratic risk implies that human capital resembles a riskless asset even more than in the baseline case. This has two effects. On the one hand, it reduces the need for precautionary savings. On the other hand, it encourages risk-taking with non-human wealth. In terms of Figure 1, one can think of individual demand functions shifting to the right. The regions of the state space where household do not save or where there is high leverage plus non-participation in equity thus become more relevant and more young and middle-aged households fall into these regions.

The experiment reported in Figure 8 builds on the previous experiment, but in addition it shuts down idiosyncratic risk on housing. It thus assumes counterfactually that every households owns a claim in a nationwide real estate investment trust. For younger households, the further reduction in risk works similarly to the reduction of non-asset income risk. In particular, savings and equity holdings drop, while borrowing increases. For older households, savings is essentially unchanged. However, there are now dramatic changes in the portfolio of the elderly, in sharp contrast to the previous experiment: elderly households now borrow a lot and invest in equity and, especially, housing.

It is natural that housing becomes more attractive as it becomes less risky. The response of the elderly households, who are not close to the borrowing constraint, can be understood as optimal rebalancing towards an asset that now has a much higher Sharpe ratio. It makes sense to go short the only asset where short positions are allowed – namely bonds – in order to exploit that Sharpe ratio. However, a reduction in housing risk also interacts with the collateral constraint. It shifts asset demand functions to the right and thus lowers savings and increases risk taking with non-human wealth. Overall, we conclude that idiosyncratic shocks play an important role in portfolio decisions. Removing risk on either type of idiosyncratic shock makes housing more attractive. This is both because of the direct effect on the housing return, and because of the role of housing as collateral.

VI Heterogeneity and Equilibrium Prices

In this section, we move beyond the analysis of household demand and consider the effect of heterogeneity on equilibrium prices. We have seen in the previous section that changes in the distribution of income and asset endowments – especially across cohorts – as well as changes in idiosyncratic risk can have large effects on asset demand. This does not mean, however, that changes in these inputs need to matter a lot for asset prices. If asset demand is very sensitive to asset prices, then it might be that the changes we have considered require only small price adjustments to get back to market clearing. Moreover, initial asset wealth – which we have taken as given in the analysis so far – is now endogenous variable. The experiments in this section change the distribution of asset endowments (as well as non-asset income). However, since prices are now allowed to change, the effect on the distribution of initial asset wealth is endogenous.

In Table 4, we collect asset prices and aggregate household sector equilibrium portfolio weights for the baseline model as well as three hypothetical combinations of inputs. In the exercise labelled "No

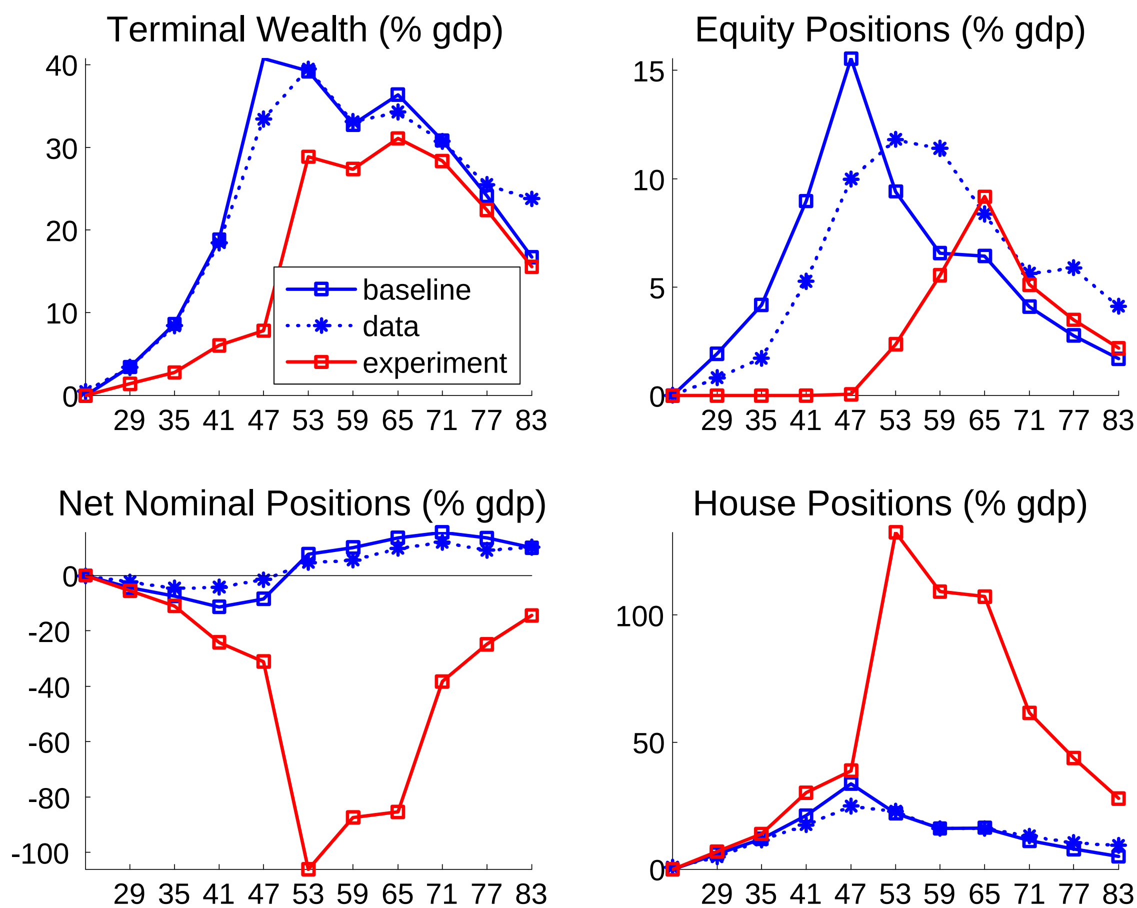

Figure 8: Asset positions by cohort in the data, implied by the baseline model, and implied by the model without idiosyncratic income or housing risk ("experiment"). The positions are in percent of GDP.

heterogeneity within age cohorts", there is a single household per cohort who is assigned all income and asset endowments owned by the cohort in the baseline case. Asset demand at fixed prices for this case was discussed in subsection B.. At the old prices, households preferred to hold less bonds – or even borrow more – but overall save more in stocks and houses than in the baseline case. In equilibrium, this leads to an 8% increase in stock prices and a 3% increase in house prices, as well as a 20 basis point increase in the nominal interest rate. There is also an overall increase in borrowing and lending among households.

The exercise labelled "All wealth owned by 47-53 cohort" is designed to provide a stark illustration of the large effects of changes in wealth and income across cohorts on prices. Here we assume that there is a single household, who is in the 47-53 year old cohort – that is, his asset demand function is that of the 47-53 year old cohort – but who is assigned all income and all asset endowment available in the economy. The effect on the prices of long-lived assets is large: price dividend ratios on stocks and houses almost double, and drive a large increase in the ratio of overall housing net worth relative to GDP. At the same time, the riskless interest rate drops by about one percent. Since inflation expectations have been kept fixed, this corresponds to a one percent drop in the real interest rate, which is also a large change by historical standards.

The exercise labelled "No heterogeneity within age cohorts, no idiosyncratic income shocks" assumes that there is a single household within each cohort who not only starts out with all asset endowment and income owned by the cohort in the baseline case, but also perceives no idiosyncratic shocks to non-asset income. Broadly speaking, the result is the opposite of the previous experiment: the lower desire to save increases the interest rate and lowers stock and house prices. Here the effect on house prices is relatively stronger than that on stock prices. This is because houses serve as collateral. Intuitively, prices must adjust to discourage the buildup of leveraged portfolios. The increase in the interest rate and a drop in the price of housing both serve this purpose.

Table 4: Asset Prices For Different Experiments

	NW/GDP	portfolio weights			credit/GDP		pd ratios		nominal
		bonds	housing	stocks	+	−	housing	stocks	rate (%)
Data	**2.51**	**15**	**59**	**26**	**70**	**31**	**20.4**	**23.9**	**6.1**
Model	2.51	15	60	25	70	31	20.8	22.7	6.1
No heterogeneity within age cohorts									
	2.60	15	60	25	73	35	21.5	24.6	6.3
All wealth owned by 47-53 cohort									
	4.50	9	64	27	38	0	40.0	45.2	5.4
No heterogeneity within age cohorts, no idiosyncratic income shocks									
	1.20	32	46	22	45	09	7.6	10.0	8.6

References

Alvarez, Fernando and Urban Jermann 2001. "Quantitative Asset Pricing Implications of Endogenous Solvency Constraints." *Review of Financial Studies* 14(4), pp. 117-151.

Brav, Alon, George M. Constantinides and Christopher C. Geczy 2002. "Asset Pricing with Heterogeneous Consumers and Limited Participation: Empirical Evidence." *Journal of Political Economy* 110, pp. 793-824.

Caballero, Ricardo J. 2006. "On the Macroeconomics of Asset Shortages." *NBER Working Paper* 12753.

Constantinides, George M., John B. Donaldson and Rajnish Mehra 2002. "Junior Can't Borrow: A New Perspective on the Equity Premium Puzzle." *Quarterly Journal of Economics* 117, pp. 269-296.

Constantinides, George M. and Darrell Duffie 1996. "Asset Pricing with Heterogeneous Consumers." *Journal of Political Economy* 104(2), pp. 219-40.

Curcuru, Stephanie, John Heaton, Deborah J. Lucas and Damien Moore 2004. "Heterogeneity and Portfolio Choice: Theory and Evidence." Forthcoming in the *Handbook of Financial Econometrics.*

DeNardi Mariacristina 2004. "Wealth Inequality and Intergenerational Links." *Review of Economic Studies* 71, pp. 743-768.

Doepke, Matthias and Martin Schneider 2006. "The Real Effects of Inflation through the Redistribution of Nominal Wealth." Working paper, UCLA & NYU.

Fernandez-Villaverde, Jesus and Dirk Krueger 2005. "Consumption and Saving over the Life Cycle: How Important are Consumer Durables?" Working Paper, UPenn.

Flavin, Marjorie and Takashi Yamashita 2002. "Owner-Occupied Housing and the Composition of the Household Portfolio." *American Economic Review,* pp. 345-62.

Gomes, Francisco and Alex Michaelides 2005. "Optimal Life-Cycle Asset Allocation: Understanding the Empirical Evidence." *Journal of Finance* 60 (2), pp. 869-904.

Gourinchas, Pierre-Olivier and Jonathan Parker 2002."Consumption over the Life Cycle." *Econometrica* 70, pp. 47-89.

Grandmont, Jean-Michel 1977. "Temporary General Equilibrium." *Econometrica* 45, pp. 535-572.

Grandmont, Jean-Michel 1982. "Temporary General Equilibrium Theory." Chapter 19 in K.J. Arrow and M.D. Intriligator, Eds, *Handbook of Mathematical Economics*, North Holland, Amsterdam.

Heathcote, Jonathan, Kjetil Storesletten and Giovanni L. Violante 2004. "The Macroeconomic Implications of Rising Wage Inequality in the United States." Working paper, NYU.

Heaton, John and Deborah J. Lucas 1996. "Evaluating the Effects of Incomplete Markets on Risk Sharing and Asset Pricing." *Journal of Political Economy* 104, pp. 443-487.

Heaton, John and Deborah J. Lucas 2000. "Portfolio Choice and Asset Prices: The Importance of Entrepreneurial Risk." *Journal of Finance* 55, pp. 1163-1198.

Hurst, Erik, Ming-Ching Luoh and Frank Stafford 1998. "Wealth Dynamics of American Families: 1984-1994. *Brookings Papers on Economic Activity* 1.

Jagannathan, Ravi and Narayana Kocherlakota 1996. "Why Should Older People Invest Less in Stocks Than Younger People?" *Quarterly Review* Federal Reserve Bank of Minneapolis 20(3), pp. 11-23

Krueger, Dirk and Felix Kubler 2004. "Computing Equilibrium in OLG Models with Stochastic Production." *Journal of Economic Dynamics and Control* 28(7), pp. 1411-1436.

Krueger, Dirk and Hanno Lustig 2006. "The Irrelevance of Market Incompleteness for the Price of Aggregate Risk'." Working paper, UPenn and UCLA.

Krusell, Per and Anthony Smith 1997. "Income and Wealth Heterogeneity, Portfolio Choice, and Equilibrium Asset Returns." *Macroeconomic Dynamics* 1, pp. 387-422.

Krusell, Per and Anthony Smith 1998. "Income and Wealth Heterogeneity in the Macroeconomy." *Journal of Political Economy* 106(5). pp. 867-896.

Krusell, Per and Anthony Smith 2006. "Quantitative Macroeconomic Models with Heterogeneous Agents." Working paper, Princeton and Yale.

Levin, David K. and William R. Zame 2002. "Does Market Incompleteness Matter?" *Econometrica* 70(5), pp. 1805-1839.

Ortalo-Magne, Francois and Sven Rady 2005. "Housing Market Dynamics: On the Contribution of Income Shocks and Credit Constraints." Forthcoming *Review of Economic Studies*.

Piazzesi, Monika and Martin Schneider 2006. "Inflation and the Price of Real Assets." Working paper, Chicago and NYU.

Piazzesi, Monika, Martin Schneider and Selale Tuzel 2007. "Housing, Consumption and Asset Pricing." *Journal of Financial Economics* 83, pp. 531-569.

Rios-Rull, Victor 1996. "Life-Cycle Economies and Aggregate Fluctuations." *The Review of Economic Studies* 63(3), pp. 465-489.

Storesletten, Kjetil, Chris I. Telmer, and Amir Yaron 2004. "Cyclical Dynamics in Idiosyncratic Labor-Market Risk." *Journal of Political Economy* 112, pp. 695-717.

Telmer, Chris I. 1993. "Asset-Pricing Puzzles and Incomplete Markets." *Journal of Finance* 48, pp. 1803-1833.

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