

INFLATION AND THE PRICE OF REAL ASSETS*

Monika Piazzesi Martin Schneider
University of Chicago NYU

July 2005

Abstract

This paper considers price and quantity movements in the three major asset classes – real estate, equity, and nominal fixed income – in the postwar period. To understand these movements, we compute a sequence of temporary equilibria in a lifecycle model with heterogeneous agents and uninsurable nominal risk. A key input to the model is the joint distribution of asset endowments and income, which we take directly from household level data. We show that changes in inflation expectations, together with demographic shifts and changes in asset supply, help understand the experience of the 1970s.

*PRELIMINARY & INCOMPLETE. Comments welcome! Email addresses: piazzesi@uchicago.edu, martin.schneider@nyu.edu. We thank John Heaton, Toby Moskowitz and conference participants at the Federal Reserve Bank of Atlanta 2005 (“Housing and the Macroeconomy”), the Berlin conference 2005 (“Macroeconomic Risk and Policy Consequences”), and seminar participants at the University of Chicago for comments.

I Introduction

Both the size and composition of US household sector wealth have changed dramatically over the postwar period. Figure 1 plots aggregate household wealth in the three major asset classes – real estate, equity, and nominal instruments, all as multiples of GDP. In both the 1960s and the 1990s, household wealth was high relative to GDP. In addition, the aggregate portfolio was relatively more tilted towards equity, as shown in Figure 2. In the 1970s and early 1980s – a period notable for relatively high inflation – wealth was lower relative to GDP, and to a greater extent invested in real estate. At the same time, the household sector’s *net* position in nominal instruments – that is, nominal assets minus nominal liabilities – has been relatively stable over the whole period, at least compared to positions in equity and real estate.

This paper develops a model of the joint evolution of asset prices and household wealth portfolios. The key features of the model are household heterogeneity by age and wealth and the uninsurability of nominal risk. We make the distribution of income and initial endowments of agents observable by measuring them based on successive cross sections of the Survey of Consumer Finances and aggregates from the Flow of Funds Accounts. We show that the model is able to capture a number of stylized facts about individual portfolio choices that we can observe in the SCF. We then use it to assess quantitatively whether changes in inflation expectations are responsible for Figures 1 and 2, and how their effect compares with demographic shifts and movements in asset supply. We find that a combination of these factors offers the most plausible account.

Our approach is based on two building blocks. The first is a model of lifecycle savings and portfolio choice at the individual level. Households choose between three assets equity, real estate and nominal bonds. They experience idiosyncratic labor income risk that cannot be insured using only the three available assets. In credit markets, they face collateral constraints – all borrowing must be backed by real estate – as well as a spread between borrowing and lending rates. Moreover, lending and borrowing are required to be in nominal terms – there is no riskfree asset. Inflation risk is thus also uninsurable.

The second building block is a one-period model of asset trading. Households of different age enter a period with different endowments of the three assets as well as non-asset income.




Figure 1: Aggregate wealth components (market value) divided by GDP, Flow of Funds & own computations, 1952:1-2003:4.

During the period they trade in asset and goods markets. Expectations about future returns and non-asset income are given exogenously. Prices clear markets in the current trading period to establish a “temporary equilibrium”, as in Grandmont (1977, 1982). In the asset and goods markets, households also interact with the rest of the economy, that is, the government, business and foreign sectors. In particular, the model allows all assets to be in nonzero net supply; supply can change within a period as the household sector trades with other sectors.

We implement the model quantitatively for three 6-year periods, centered around benchmark years 1968, 1978 and 1995. The idea here is to view Figure 2 as a sequence of temporary equilibria, of which three are explicitly computed. For each period, we measure (i) the joint distribution of initial asset endowments and income endowments, (ii) the supply of assets to the household sector by other sectors. Data for these model inputs come from the Flow of Funds accounts and the Survey of Consumer Finances. The other key model input is households’ expectations about future asset returns and income. Here we determine a set of baseline expectations from empirical moments.

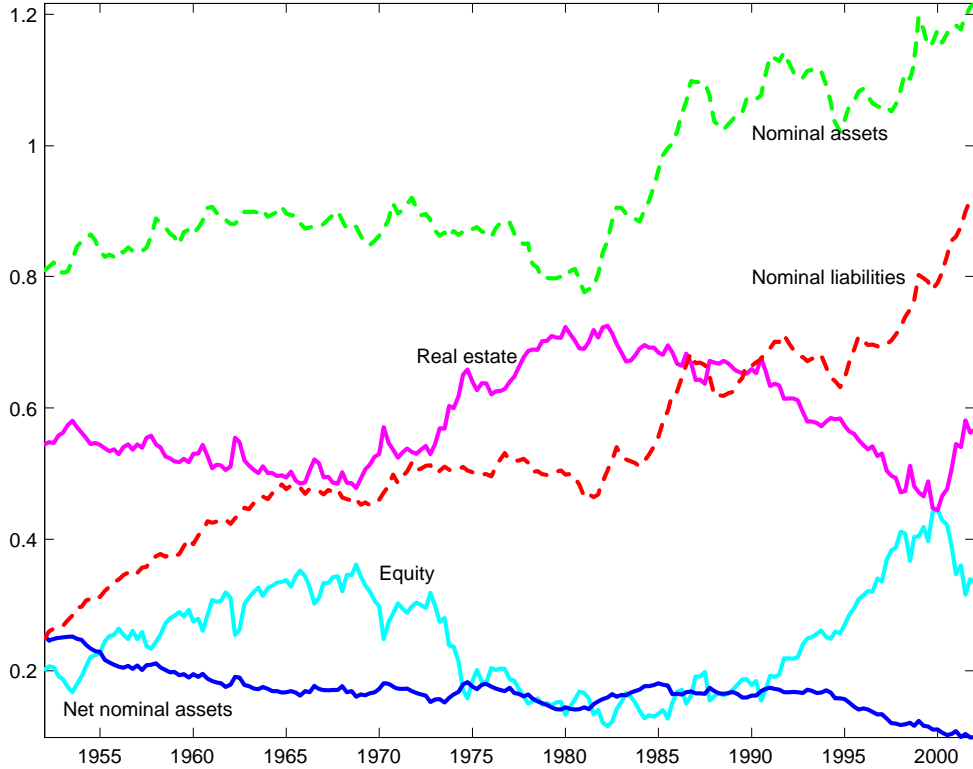


Figure 2: Aggregate portfolio shares, Flow of Funds & own computations, 1952:1-2003:4.

We calibrate preferences so that, under baseline expectations, the model replicates the aggregate wealth-to-GDP ratio as well as household sector portfolio weights for 1995. This requires a coefficient of relative risk aversion of 15 and a discount factor slightly above one. We show that, at these parameter values, the model also captures a number of stylized facts about the cross-section of households in the 1995 SCF. In particular, it generates realistic 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.

We then use the model to examine the 1970s. We first show that the model under baseline expectations replicates (i) portfolio shares, prices and interest rates for 1968 and (ii) the drop in aggregate wealth between 1968 and 1978. This result follows from changes in demographics and bond supply. However, the baseline model cannot account for the portfolio shift towards housing. We thus explore whether changes in inflation expectations, for which we obtain three results.

First, an increase in expected inflation is consistent with a portfolio shift from stocks to houses that leaves the share of bonds unchanged. However, when the distribution of inflation expectations is matched to the Michigan inflation survey, the resulting portfolio shift is too small. Second, an increase in the conditional volatility of nominal bond returns generates a portfolio shift, and also helps explain the slowdown in gross nominal borrowing and lending in the 1970s. However, if it were to explain the portfolio shift by itself, it would lead to a nominal interest rate that is too high. Finally, changes in expected stock returns alone generate large portfolio shifts, but small changes in the nominal interest rate. Mild pessimism about stocks can thus account for the observed for the observed shift without upsetting other desirable features of the model. However, pessimism alone leads to a nominal interest rate that is too low. We thus conclude that a combination of the three feature is most likely responsible for the shift.

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 presents results under baseline expectations, both for aggregates and for the cross section of holdings. Section V considers the effect of inflation. Section VI concludes.

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 uninsurable idiosyncratic income shocks. Households can invest in three types of assets: long-lived equity and real estate as well as short lived nominal bonds. There is no riskless asset and 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 differences in planning horizon: the idea is that all agents expect to reach a certain age. We now describe the problem of a typical consumer with a planning horizon of $T > 0$ periods beyond the current period t .

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 $U_T = C_T$. 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.

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 cum-dividend price \tilde{p}_t^e ; they cannot be sold short. A tree pays d_t^e units of dividend during period t *after* trading in the equity market has taken place. The ex-dividend price is denoted p_t^e . 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 $\tau \geq t$.

To make our timing conventions for period t more concrete, consider a consumer who enters period t with an endowment $\bar{\theta}_t^e$. He can sell his shares for $\tilde{p}_t^e \bar{\theta}_t^e$ or buy new shares at the price \tilde{p}_t^e . Suppose trading in the stock market leaves the consumer with θ_t^e shares. He immediately receives $d_t^e \bar{\theta}_t^e$ units of numeraire as dividend, which can be consumed (or sold in the numeraire good market) in t . In addition, the consumer expects to receive $R_{t+1}^e p_t^e \theta_t^e$ units of numeraire once he sells his shares in period $t + 1$.

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 cum-dividend price \tilde{p}_t^h during period t ; 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 in period t . As with equity, this dividend arrives *after* trading in the real estate market has occurred. Moreover, every house requires a maintenance cost of m units of numeraire, also after trading.

Housing services can be either consumed or sold. There is a perfect rental market, where housing services can be rented at the price p_t^s . If a consumer buys θ_t^h units of real estate, he obtains a dividend $(p_t^s - m)\theta_t^h =: d_t^h \theta_t^h$ and the ex-dividend value of his property is $(\tilde{p}_t^h - d_t^h) \theta_t^h =: p_t^h \theta_t^h$. Alternatively, a houseowner may directly consume the housing services provided by his house. In either case, we assume that 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 , 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.¹ 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.² 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 $\{R_\tau^b\}$. They believe that $R_\tau^b = 1/q_\tau\pi_{\tau+1}$ is the (ex post) real lending rate, and that $R_\tau^b(1 + \xi)$ 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.

Budget Set

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

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

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

¹To 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 \$.

²One 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 .

The plan a_t must 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 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 $a_\tau = \{c_\tau, s_\tau, \theta_\tau^h, w_\tau^h, w_\tau^e, w_\tau^{b+}, w_\tau^{b-}\}$ subject to

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

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 (3)-(5).

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 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. Their precise form will be discuss in the calibration section below.

Oldest 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 the period t goods markets to maximize (1) subject to the budget constraint

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

B. Equilibrium

Suppose that there is a finite number of consumers, 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 $T(i)$.

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

$$(6) \quad C_t^{ROE} = \tilde{p}_t^h f_t^h + \tilde{p}_t^e f_t^e + D_t - \bar{B}_t.$$

If C_t^{ROE} is positive, it is consumed by the ROE sector. More generally, the ROE sector 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:

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

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 , $(\tilde{p}_t^h, \tilde{p}_t^e, q_t, p_t^s)$, 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 $a(i) = (a_t(i), \{a_\tau(i)\}_{\tau=t+1})$ for that maximizes utility 2 given consumer i 's endowment, planning horizon, and expectations about future prices and returns;

(2) markets for all assets and goods clear:

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

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 $\tilde{p}_t^h, \tilde{p}_t^e, q_t$ and p_t^s .

C. Discussion of the Assumptions

Temporary Equilibrium

In our setup, we do not require impose structural knowledge – that is we do not assume that there is a “true” probability distribution of fundamentals and a “true” model that all agents know and use for forecasting. In particular, household expectations are not required to anticipate structural change that takes place later. This reflects our view that many of changes that are relevant for thinking about asset markets – such as different monetary policy regimes – are hard to foresee, or, more generally, hard to form probabilities about. The rational expectations assumption seems more suited for models where fundamentals change in regular ways and where agents have plenty of data and experience about the link between fundamentals and prices.

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.

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.

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 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. To adjust the amount of housing, agents have to pay adjustment costs. In our model, we abstract from these issues (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.) These considerations suggest that our model may correctly capture the tradeoffs faced by these households, after we calibrate average housing returns to reflect adjustment costs.

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. This section describes where we obtain these inputs. We start with a description of what a model period corresponds to in the data.

Timing

The length of a period is six years. 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 9 age groups of households (<29, 30-35, 36-41, 42-47, 48-53, 54-59, 60-65, 66-71, >72) 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 divide the period 1959-2003 into 9 six-year periods, 1959-64, 1965-70, 1970-75, 1975-80, 1980-85, 1986-91, 1992-97 and 1998-2003. In our exercises below, we calculate temporary equilibria for three of these periods, namely 1965-70, 1975-80 and 1992-95. To construct asset endowment distributions, we use data on asset holdings from the respective precursor periods. Moreover, we pick the income distribution by specifying a stochastic process for individual income, and then forecasting the whole income distribution one period ahead. This approach allows to capture the correlation between income and *initial* asset holdings that is implied by the joint distribution of income and wealth.

Since the model compresses what happens over a six year span into a single date, prices and holdings are best thought of as averages over the period. However, individual level 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 one in the 4th year of the period, in particular, the surveys we use are 1962, 1983, 1989, 1995 and 2001 (some periods also contain a second survey in their 1st year). We use the 4th year survey to infer income and asset holdings where possible.

A. Data

The three assets in the model correspond to three broad asset classes. For each class, we identify individual positions from the SCF and aggregates from the Flow of Funds Accounts (FFA) and NIPA. Equity is identified with shares in corporations held by households. Here we include both publicly traded and closely held shares, both foreign and domestic equity, and both shares held directly by households and shares held indirectly through investment intermediaries such as mutual funds or defined contribution pensions plans. Individual positions come from the SCF of selected years. An annual series for the aggregate value is constructed from the FFA, with corresponding dividends constructed from NIPA. Properties of the resulting return series are described below. Our concept of equity does not include noncorporate business and equity held in defined benefit pension plans. Income from both sources are treated as part of non-asset income.

Our concept of residential real estate contains owner-occupied housing, directly held residential investment real estate, as well as residential real estate held indirectly by households through noncorporate businesses. An aggregate measure of the latter is available in the FFA. To construct a measure at the individual level, we multiply the household's value of noncorporate business interests by the aggregate ratio. This simple way of accounting for indirect holdings of residential real estate allows us to match almost all aggregate holdings, since very few residential properties are owned by corporations. We take aggregate dividends to be housing consumption net of maintenance from NIPA. The annual aggregate housing wealth series gives rise to a return series.

Our concept for a household's bond holdings is their net nominal position, that is, the market value of all nominal assets minus the market value of nominal liabilities. As for stocks, asset holdings include not only direct holdings, but also indirect holdings through investment intermediaries. For example, when households own a mutual fund, an estimate of the part of the fund invested in stocks is added to stock holdings, and similarly for bonds. Similar calculations are performed for defined contribution pension plans. To calculate market value, we use the market value adjustment factors for nominal positions in the U.S. from Doepke and Schneider (2004).

By non-asset income we mean the sum of labor income, transfers from the government, payoff from defined-benefit pension plans, as well as dividends from noncorporate business except those from indirectly owned residential investment 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. We subtract personal income tax on non-asset income. In particular, 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.

For the years in which both data sets are available, we check consistency of the aggregates from the SCF and FFA. As is known from the work of Antoniewicz (2001), the match is good for 1989 and 1995. However, our 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 is only about 50% of the FFA 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 that

the FFA aggregates are correct throughout, and that individual positions in the SCF suffer from proportional measurement error. We then multiply each individual position by the ratio of the FFA aggregate and the SCF aggregate for the same asset class.

B. Measuring the Distribution of Assets 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 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 periods after 1983, 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 first extract the top 10% of households by net worth. Among the bottom 90% by net worth, we then divide by homeowner and renter. Among homeowners, we divide households further into “borrowers” and “lenders.” Here a household is a borrower if his net nominal position – nominal asset minus nominal liabilities – is negative.

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 most properties that we are interested in are retained. In particular, because there are very few renters among the top 10% per

age group, the fraction of renters is very close in the two distributions. In addition, sorting into borrowers and lenders, together with the fact that very few among the top 10% are net nominal borrowers, implies that the gross amount of credit is also very close. Finally, net nominal positions as well as house and stock positions – both by age and when net worth is split as top 10% and the rest, our key measures of concentration described earlier – are directly retained.

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

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

Updating Nominal Positions

For the household’s nominal assets, the above approach does not work since these assets are short-term in our model. Instead, we determine the market value of a household’s 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

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

Letting g_t denote real GDP growth and \bar{D}_t the aggregate net nominal position as a fraction of GDP, we have

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

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. 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 3% per year spread for the borrowing rate.

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. Here we use the same process that agents in the model use to forecast their non-asset income and is the 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 can be used to determine the initial endowment and income distribution. The only difficulty here is the adjustment of exiting and entering households. Here 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 1973 aggregates from the FFA, *(ii)* the share of an individual cell member's holdings in the aggregate holdings is the same as in 1962 and *(iii)* the share of all cell members' holdings in their respective cohort aggregates is the same as in 1962.

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. Condition *(iii)* imposes 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 1973 aggregates – especially gross nominal assets – appear more similar to 1962 than to the 1980s. Once we have a distribution of positions at the cell level for 1973, we proceed as above to generate an updated distribution for 1978.

C. Distributions for 1968, 1978 and 1995

Figure 3 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,

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

This aggregate will be useful to interpret the results below.

Two demographic changes are apparent from the figure. First, the baby boom makes the youngest cohort the largest group in the 1978 cross section. By 1995, the boomers have aged so that the 42-47 year olds are the now strongest 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 for the oldest group 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 elderly as well as the youngest. While the 1995 distribution also has relatively more weight on the elderly, it emphasizes more the middle-aged rather than the young.

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.

D. 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 top panel of Figure 4 plots both quarterly series of these numbers. The calibration of the model uses six-year aggregates.

The initial nominal positions 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 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 FFA for period t . This series is reproduced in

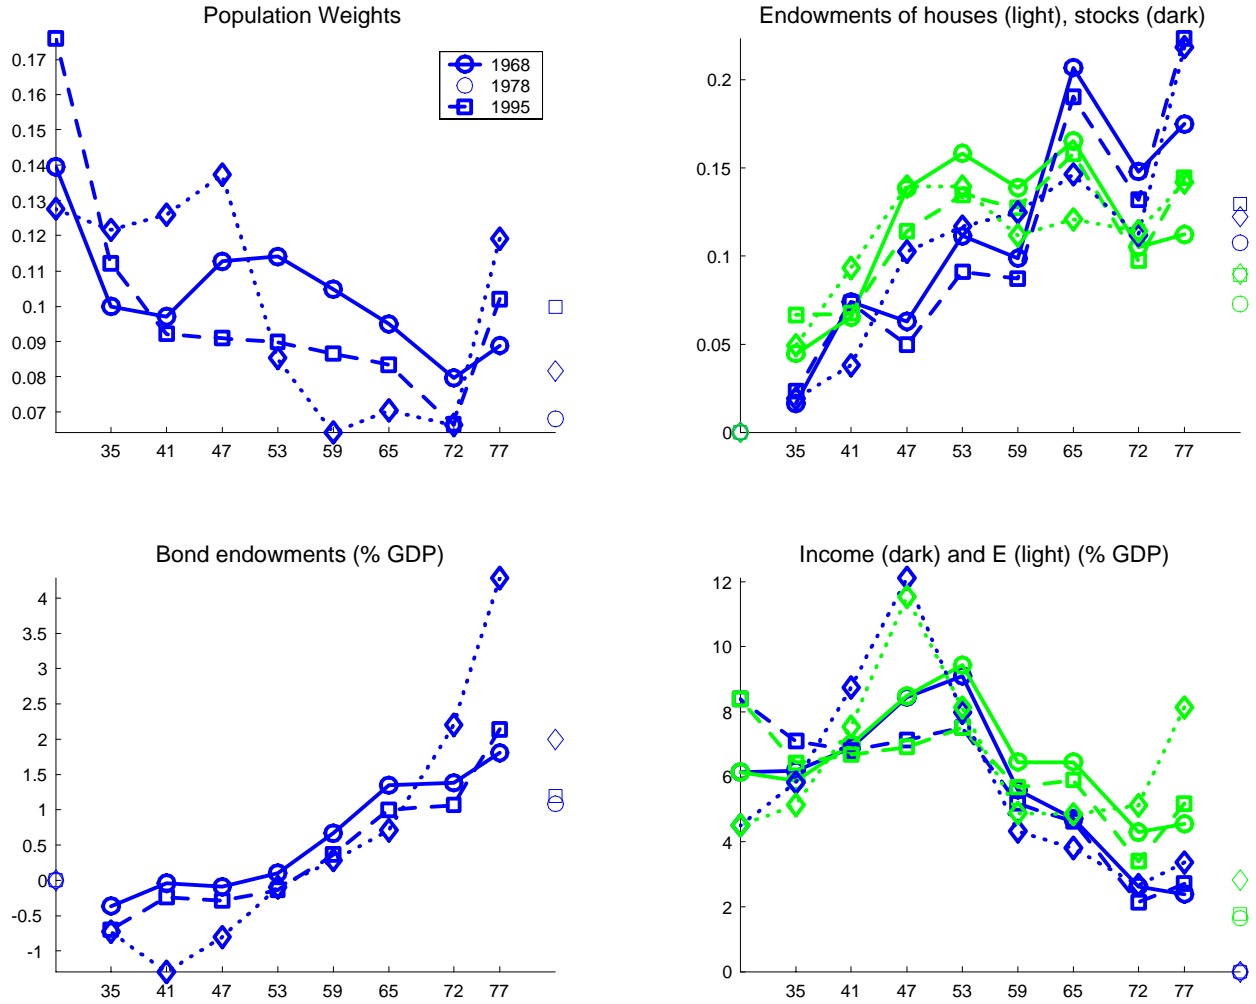


Figure 3: Asset endowment and income distributions in 1968, 1978 and 1995. *Top left panel:* Population weights by cohort, identified on the horizontal axis by the upper bound of the age range. Exiting households during the period are on the far right. *Top right panel:* House endowments (light lines) 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 bottom panel of Figure 4.

E. Baseline expectations

The previous subsections have described actual income and holdings in a trading period t . In addition, the model requires agents' expectations about returns and income in the future.

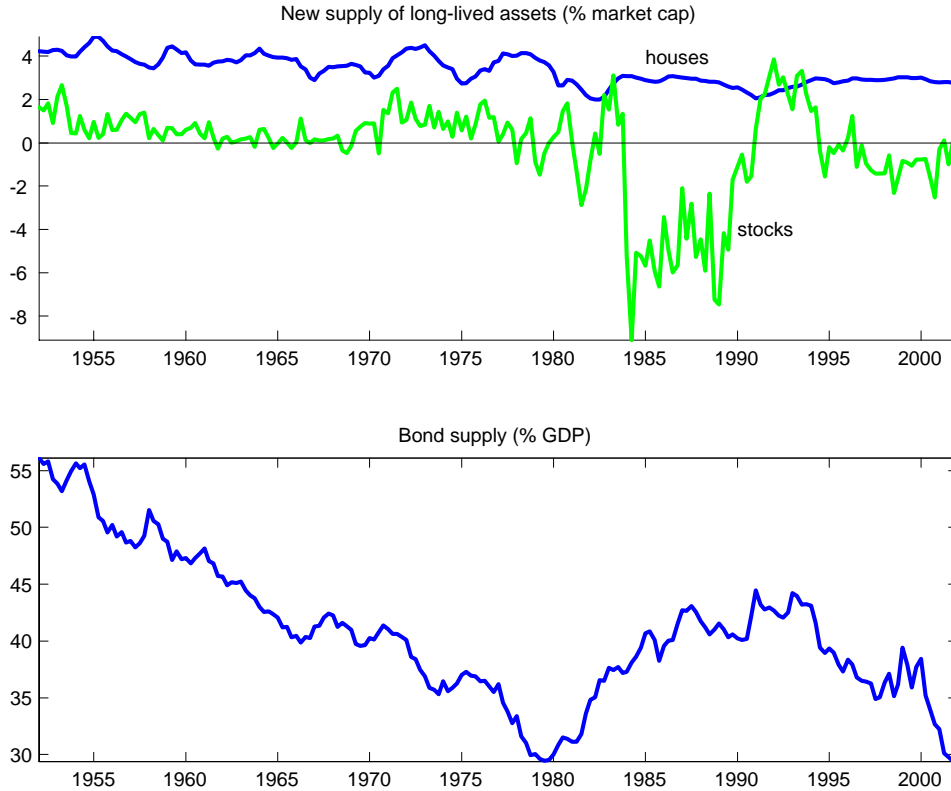


Figure 4: *Top panel:* Net new corporate equity as a percent of total household holdings of corporate equity and residential investment as percent of value of residential real estate. Data are quarterly at annual rates. *Bottom panel:* Net nominal position of household sector as a percent of GDP.

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 (2001), 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.

Estimates of labor income processes at the annual level tend to produce estimates for σ_ε between 1% and 2% per year. We adopt $\sigma_\varepsilon = .02$ here. For σ_u , estimates on annual data yield estimates of σ_u that suggest that temporary shocks are about 10 times more volatile than temporary shocks. However, this result holds at the annual level – over six year periods, some of the iid temporary shocks will wash out. We proceed with $\sigma_u = .1$. Importantly, we use these numbers only for volatility during “working life”, which we take to be age 65 and younger. There is no idiosyncratic uncertainty after retirement. Effectively, we are assuming that the consumer obtains a sure indexed pension.

TABLE 2: INCOME AGE PROFILE

	35	41	47	53	59	65	71	77+
1962	1.27	1.42	1.49	1.48	1.36	0.97	0.73	0.61
1995	1.41	1.78	2.10	2.51	2.06	1.65	1.04	0.70

NOTE: Income age profiles estimated from the 1962 and 1995 SCFs. The numbers represent the average cohort income relative to the average income of the youngest cohort (≤ 29 years).

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 2 reports the profile relative to the income of the youngest cohort.

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 3 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 flow of funds accounts (FFA), we construct returns on corporate equity and residential real estate directly from FFA aggregates.

TABLE 3: SUMMARY STATISTICS

r_b	r_h	r_s	π	g
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 6-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 is computed using the CPI, while g 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 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 National Income and Products Accounts (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 3 directly for our benchmark beliefs. However, this would not capture the tradeoff faced by the typical individual household. Indeed, the housing returns in Table 3 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 3 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 variance 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 et al. (2005) 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.³

³Since the volatility of housing is measured imprecisely, we chose the precise number for the multiplicative factor

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.

IV Supply, Demographics and Asset Prices

In this section, we compute equilibria at baseline expectations for 1968, 1978 and 1995. This isolates the effect of changes in supply and the wealth distribution on asset prices. We then compare the performance of the model in the cross-section of households to actual observations from the 1995 SCF.

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.

Baseline Parameters

The previous section delivers distributions of asset endowments and income for each year. It remains to choose preference and credit market parameters. We fix an intertemporal elasticity of substitution of $\sigma = .5$, a standard value in the literature, and let $\delta = .86$, the share of housing services in aggregate consumption in the data. This expenditure share does not vary much over the lifecycle and across households (Piazzesi et al. 2005, Appendix B). We assume that there is a 3% per year spread between borrowing and lending interest rates. In addition, we select the borrowing constraint parameter $\alpha = .8$. This implies a maximal loan-to-value ratio of 80%, where “value” is the ex-dividend value of the house.

Finally, we select a coefficient of relative risk aversion of $\gamma = 15$ and a discount factor $\beta = e^{-.06}$. At these values, the model roughly matches aggregate portfolio weights and the aggregate wealth-to-GDP ratio in 1995. The numbers for the model and the data are reported in Table 4. High risk aversion is required to entice agents place a large enough weight on bonds in equilibrium. A high discount factor – in fact, a negative rate of time preference – is required for the model to generate sufficient savings out of initial wealth.

Baseline Results on Aggregates

Comparison of panels A and B of Table 4 shows that the model does a decent job not only for key aggregates in 1995, but also for 1968. The model thus captures the fact that household portfolios in these two years are quite similar: total wealth is roughly 2.5 times GDP, with about half invested in real estate and about 30% in stocks. The main differences between the years are that the wealth-to-GDP ratio and the portfolio share on houses rise from 1968 to 1995, while the portfolio shares on bonds and stocks fall. The model generates the same qualitative changes for bonds, although the portfolio shift out of stocks into real estate and the increase in wealth are smaller.

In terms of asset prices, the price-dividend ratios for housing and stocks were similar in the years 1968 and 1995. The model does a pretty good job in capturing these ratios. The level of interest rates was also similar, at about 6%. The model predicts interest rates of about 7% for both years. For the data, we report numbers for gross credit from the previous survey for 1968 (because

we do not have a survey from the same year), constructed numbers from the updating mechanism discussed previously, and the actual 1995 SCF numbers. While the 1978 number is very tentative, we can see that gross credit increased over the years.

TABLE 4: BASELINE RESULTS

year	beliefs	portfolio weights			lend./ GDP	borr./ GDP	wealth/ GDP	PD ratios		interest rate nominal
		bonds	housing	stocks				housing	stocks	
PANEL A: DATA										
1968		.17	.49	.34	.64	.18	2.47	20.2	32.6	.057
1978		.15	.69	.16	.59	.23	2.03	25.0	14.5	.084
1995		.15	.56	.29	.71	.32	2.63	20.1	28.1	.061
PANEL B: MODEL										
1968	baseline	.17	.53	.30	.56	.15	2.44	21.5	30.1	.071
1978	baseline	.15	.53	.32	.46	.15	2.02	19.2	34.7	.072
1995	baseline	.15	.55	.30	.62	.25	2.50	18.8	27.1	.070

NOTE: Panel A reports the aggregate portfolio weights on bonds, housing and stocks from Figure 2; the gross borrowing and lending numbers from 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 1978: 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 model remains essentially the same as in the two other years. As a result, the price dividend ratios of houses falls and that of houses rises, in contrast to what happened in the data.

Two changes in fundamentals are important for the drop in the wealth-to-GDP ratio in 1978. The first is the demographic effect illustrated in Figure ?? of Section C. 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 will be shown in the next section, the model predicts that these cohorts have small savings rates. This leads to lower wealth-to-income ratios and pushes interest rates up. However, the latter effect is counteracted by the second important change – the 25% reduction in bond supply documented in Figure 4. Taken together, these effects produce relatively stable interest rates, lower wealth-to-GDP ratios and dramatic shifts in the composition of the aggregate portfolio.

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 – that is asset wealth plus non-asset income – to the permanent component of non-asset income. For simplicity, we refer to the latter ratio as the wealth to income ratio. Figure 5 plots agents decisions as a function of this wealth-to-income ratio.

Savings

The bottom right panel shows the ratio 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.

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 below 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. 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, borrowing serves to construct a leveraged portfolio. 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 Cocco (2005) and Campbell et al. (2003).

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.

The model can currently not 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. Experimentation with alternative beliefs has shown that if stocks are relatively less attractive, it is possible to obtain a stock share in the portfolio that increases with wealth-to-income for intermediate levels of the latter. The behavior of the portfolio weight on stocks implies that the model produces typically too little concentration of stock ownership.

Preliminary results on a version of the model with owner-occupied and rental housing suggests that the latter feature helps along several of the dimensions where the current version is still lacking. In particular, if housing services can be consumed more cheaply when there is owner occupation, there is an additional reason for young people to hold relatively more houses. This will also contribute to making stock ownership more concentrated.

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 Survey of Consumer Finances. 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, say, houses for a household i can be written as

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

where $\bar{\alpha}^h$ is the aggregate portfolio weight on houses. A model that correctly predicts the cross section of portfolio shares will therefore only correctly predict the cross section of market shares

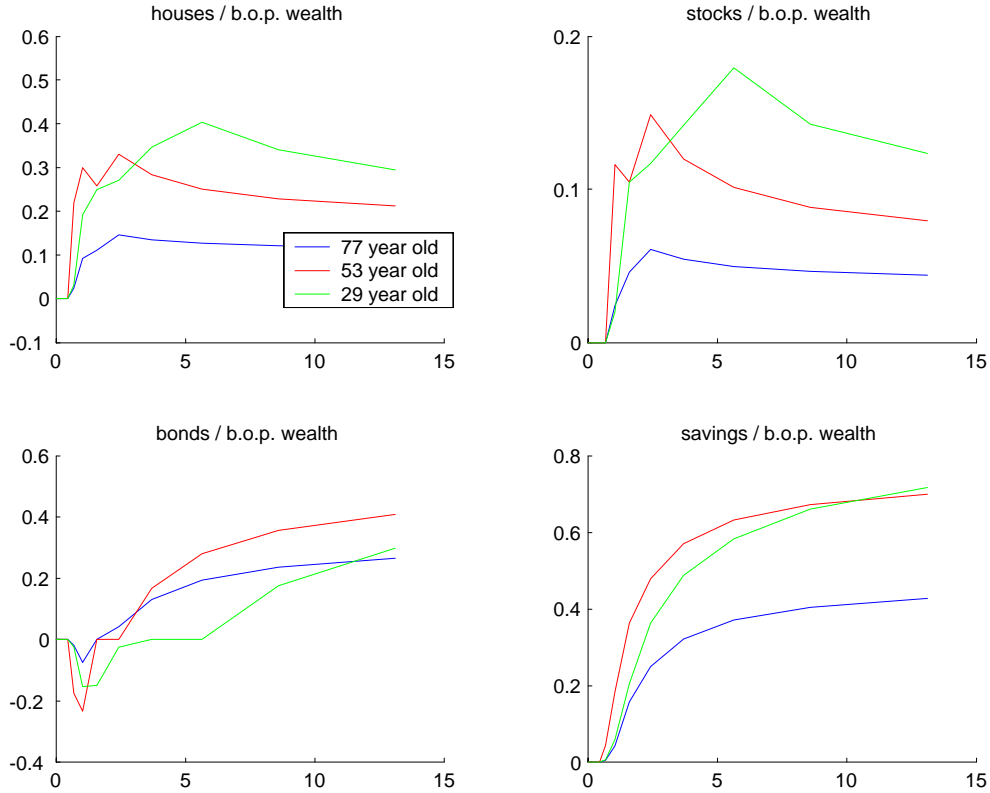


Figure 5: Asset holdings and terminal wealth, both as fractions of initial wealth, plotted against the initial wealth-to-income ratio. Age groups are identified by maximum age

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 model does a fairly good job at matching terminal wealth. One exception is the very oldest group of savers who save too little in the model. The model also captures skewness of the distribution of terminal wealth and how this skewness changes with age. The top 10% by net worth own more than half of total terminal wealth, their share 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 increase 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 generate 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

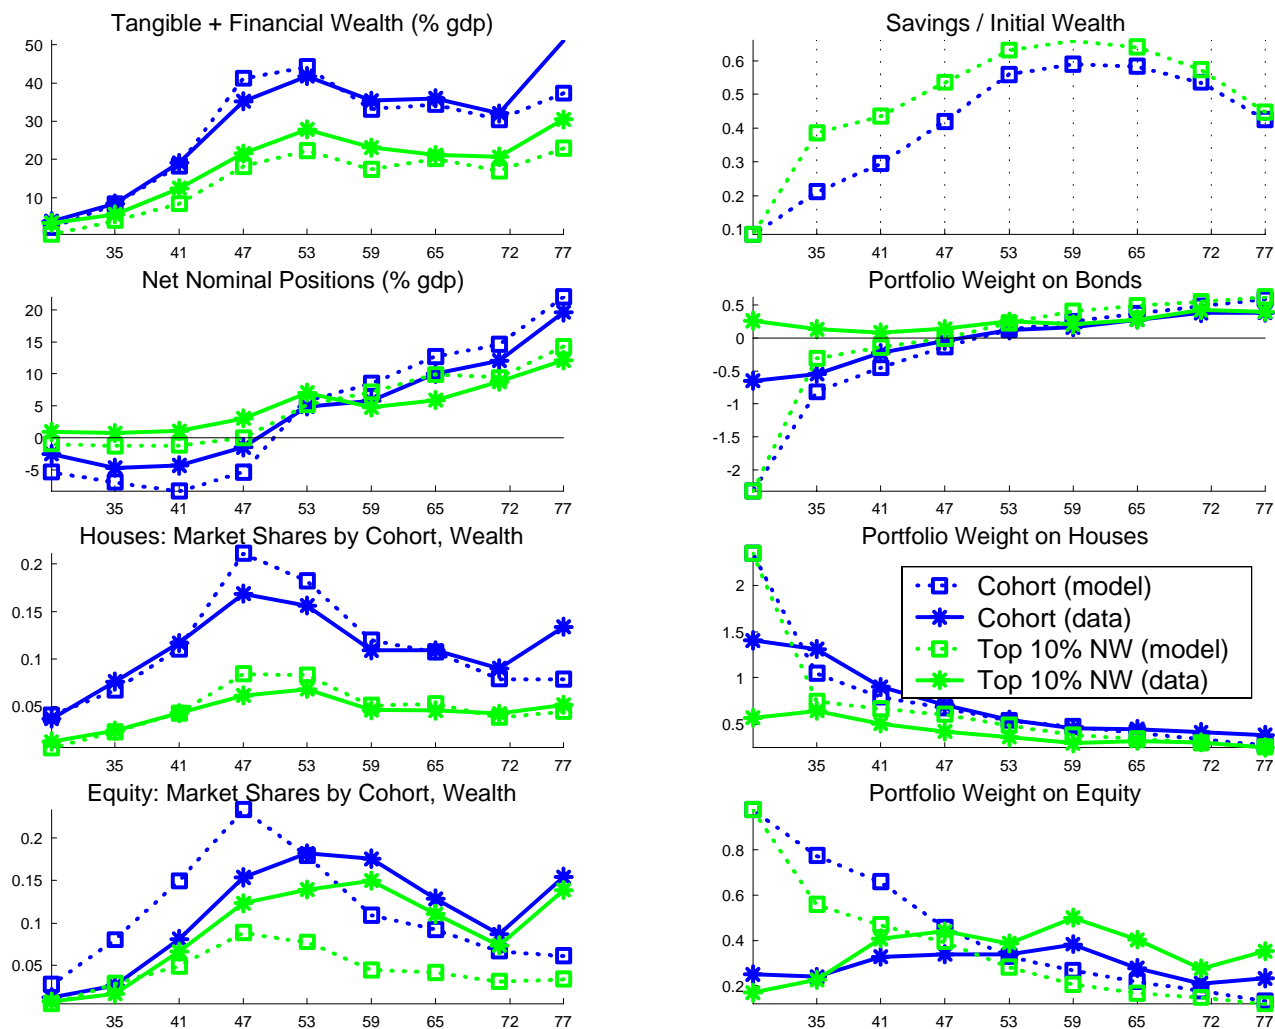


Figure 6:

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 The Effects of Inflation

In this section, we use our model to explore the effects of inflation on the price of real assets. The idea is to see whether the effects of (i) expected inflation through capital gains taxation together with (ii) inflation uncertainty and (iii) lower expected firm profitability can help us understand why stock prices fell in the 1970s while house prices rose. Key statistics for the various scenarios are reported in Table 5. To facilitate quantitative assessments and comparisons, we repeat the main facts in the first row of Table 5 (where real rates are not observable) and the “baseline” exercise based on baseline beliefs in the second row.

The case of “hi inflation expectations” increases the mean inflation rate from 4% to 7%, which is the median forecast in the Michigan survey for inflation over the next 5 years. Figure 7 plots the median forecasts, which are available starting in 1978 together with median forecasts for different age cohorts. Interestingly, median forecasts by older households during the great inflation tend to be lower than forecasts by younger households. Our lifecycle model offers a natural laboratory for studying the impact of these heterogenous expectations – the case “heterogeneous inflation expectations” in Table 5 is based on these cohort forecasts from the Michigan Survey. Finally, the case of “hi inflation volatility” increases the volatility of π_{t+1} . Since the Michigan survey does not ask for inflation volatility estimates, we set the volatility such that the model roughly matches the aggregate portfolio weight on housing. This approach leads to an increase of volatility by a factor of 8.

Table 5 illustrates that changes expected inflation and inflation uncertainty both increase interest rates and have significant – and opposite – effects on the value of stocks and houses. The effects are driven by the nonneutrality of capital gains taxes. For example, households face after-tax real returns on equity $(1 - \tau) r_t^e - \tau \pi_t$ and bonds $(1 - \tau) r_t^b - \tau \pi_t$, while the real return on housing r_t^h is untaxed. In the case of “hi inflation expectations”, capital gains taxes imply a larger tax burden $\tau \pi_t$. Since the tax burden only affects taxable assets, households will shift their portfolio towards housing, which is not taxed. As a consequence, equity and bonds will become less valuable, while housing will become more valuable. The effect on house prices is reinforced by the tax-deductibility of mortgages, since the portfolio shift away from bonds implies that households will want to borrow




Figure 7: Median inflation expectations over the next 5 years, Michigan Survey of Consumers.

and invest even more into housing. Indeed, rows 3 and 4 of Table 5 show a slight increase in gross borrowing and lending due to the increased real tax subsidy associated with mortgages. This expected inflation story was advanced by Feldstein (1982) using deterministic valuation models.

Row 4 in Table 5 considers the “heterogenous inflation expectations” case, where households’ expectations differ by age cohort as measured by the Michigan Survey. Here, older households, who tend to have higher wealth-income ratios and thus save more, anticipate inflation to be lower than younger households, who tend to be borrowers. Older households are therefore happy to lend at nominal rates that are viewed as bargain by young households. As a result, gross borrowing and gross lending goes up, which is the main difference to the case of homogeneous expectations. The other implications are similar: stocks and bonds become cheaper, while houses become more expensive.

Row 5 in Table 5 considers the “hi inflation volatility” case, where households view inflation as more uncertain. This case is unique to our model, since it requires the modeling of risk. In our

setup, inflation determines real after-tax returns on stocks and bonds. To the extent that households are more uncertain about inflation, they will view the return on taxable assets as more uncertain and, again, shift their portfolio towards housing. Again, this will increase the value of housing and lower the value of stocks and bonds. However, unlike in the case of “hi inflation expectations,” households will not want to borrow more. The reason is that leverage would increase the volatility of their overall portfolio without adding returns on average. When inflation uncertainty is high enough, the market of intra-household borrowing and lending will break down completely, as in the 5th row of Table 5.

To summarize the two stories, both higher “hi inflation expectations” and “hi inflation volatility” are able to generate the kind of the movements in real asset prices that we observe during the 1970s: house prices go up, while stock prices go down. Also, the two stories produce higher nominal interest rates, another hallmark of the experience in the 1970s. However, we find that nominal rates are highly sensitive to inflation expectations. To generate pd-ratios for stocks and housing below and above 20, respectively, the model requires nominal rates above 10%, which seems high.

In rows 6-8 of Table 5, we examine another potential explanation for the inverse relation between house and stock prices during the 1970s: lower stock return expectations. Such a scenario seems plausible, because times of higher inflation are typically times of lower firm profitability. For example, historic cost depreciation and FIFO inventory accounting increase the tax burden on corporations (Feldstein and Summers 1979). In these episodes, households will rationally expect lower returns on their stock investments. In rows 6-8, we report results for return expectations that are 2, 3 and 4 percentage points lower than the historical equity premium. As households lower their return expectations for stocks, other assets become relatively more attractive and thus valuable. Remarkably, the price-dividend ratio for stocks and housing are highly sensitive to households’ subjective equity premium, while interest rates are not. Hence, this scenario is able to generate large movements in the price of real assets, while bond returns remain relatively stable. If households lower their stock return expectations between 3 and 4 percentage points, the model is able to match the price-dividend ratios observed in the data.

While lower stock return expectations can explain the price movements in real assets, they do not explain the runup in nominal rates. Therefore, the results in rows 6-8 suggest that we will need

a combination of lower stock expectations and some change in inflation expectations to explain the 1970s. To tell apart higher inflation expectations from higher inflation uncertainty, we need data on gross borrowing and lending during the great inflation. When we compare the changes in gross credit over the years in Panel A of Table 4, we can see that cross credit went down during the 1970s. If inflation expectations were indeed heterogeneous as measured by the Michigan Survey, this suggests that an increase in inflation uncertainty must be at least part of the story.

TABLE 5: THE EFFECTS OF INFLATION

beliefs	portfolio weights			lend./ GDP	borr./ GDP	wealth/ GDP	PD ratios		interest rate	
	bonds	housing	stocks				housing	stocks	nom.	real
—1978 data—	.15	.69	.16	.59	.23	2.03	25.0	14.5	.084	
baseline	.15	.53	.32	.46	.15	2.02	19.2	34.7	.072	.032
hi infl exp	.15	.56	.29	.46	.15	2.01	19.9	31.6	.108	.038
hetero infl exp	.15	.57	.28	.61	.30	2.02	20.5	30.7	.113	
hi infl vol	.16	.67	.17	.31	.00	1.97	23.4	19.1	.168	.128
stock exp - 2%	.16	.62	.22	.48	.17	1.97	21.7	23.6	.069	.029
stock exp - 3%	.16	.67	.17	.49	.18	1.97	23.5	18.7	.066	.026
stock exp - 4%	.15	.72	.16	.49	.19	1.95	25.2	13.7	.063	.023

NOTE: Rows 1 and 2 of this table repeat the results from Table 4. Row 3 with “hi inflation expectations” reports results when households expect inflation to be 7%. Row 4 with “heterogeneous inflation expectations” reports results when different cohorts have different inflation expectations measured using the Michigan Survey of Consumers. Row 5 with “high inflation volatility” reports results when expectations are based on an 8 times higher volatility of inflation. Rows 6-8 report results based on expectations that lower mean stock returns by the amount indicated.

VI Conclusion

In this paper, we have combined aggregate data from the Flow of Funds with household-level data from successive SCF cross sections. This approach allows us to measure the income and asset endowment distribution across households at the beginning of each trading period. To explicitly capture nonstationarities, we consider a sequence of temporary equilibria of this heterogeneous agent

economy. There are three assets – housing, stocks and nominal bonds. There is no riskless asset, so that market are incomplete. During the 1970s, households anticipate higher inflation and view inflation as more uncertain. In particular, young households adjusted their inflation forecasts more than old agents. These changes in inflation expectations make housing more attractive, because of capital gains taxes on stocks and mortgage deductibility. The effects associate with these changes alone, however, cannot explain the opposite movements of house and stock prices during the great inflation, because they have dramatic impacts on interest rates. To fully explain the asset prices of real assets, we need to also allow agents to interpret inflation as bad news for future stock returns. Taken together, these effects can then explain the opposite movements of house and stock prices in the 1970s.

References

- [1] Campbell, John Y. and Tuomo Vuolteenaho. 2004. "Inflation Illusion and Stock Prices." *American Economic Review* 94, pp. 12-3.
- [2] Campbell, John Y. and Joao Cocco. 2003. "Household Risk Management and Optimal Mortgage Choice." *Quarterly Journal of Economics* 118, pp 1449-1494.
- [3] Caplin, Andrew, Sewin Chan, Charles Freeman, and Joseph Tracy 1997. "*Housing Partnerships*." Cambridge and London: MIT Press.
- [4] Cocco, Joao F. 2005 "Portfolio Choice in the Presence of Housing." *Review of Financial Studies* 18, pp. 535-567
- [5] Cocco, Joao F., Francisco Gomes and Pascal Maenhout 2005. "Consumption and Portfolio Choice over the Lifecycle." *Review of Financial Studies* 18, pp. 431-533.
- [6] 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, 269-296.
- [7] Constantinides, George M. and Darrell Duffie 1996. "Asset Pricing with Heterogeneous Consumers." *Journal of Political Economy* 104(2), pp. 219-40.
- [8] Curcucu, Stephanie, John Heaton, Deborah J. Lucas and Damien Moore 2004. "Heterogeneity and Portfolio Choice: Theory and Evidence." Forthcoming in the *Handbook of Financial Econometrics*.
- [9] Davis, Steven J., Felix Kubler and Paul Willen. 2005. "Borrowing Costs and the Demand for Equity over the Life Cycle." Working paper, Chicago GSB, Mannheim & Boston Fed.
- [10] DeNardi Mariacristina 2004. "Wealth Inequality and Intergenerational Links." *Review of Economic Studies* 71, pp. 743-768.
- [11] Doepke, Matthias and Martin Schneider 2004. "The Real Effects of Inflation through the Redistribution of Nominal Wealth." Working paper, UCLA & NYU.

- [12] Feldstein, Martin 1980. "Inflation and the Stock Market." *American Economic Review* 70(5) pp. 839-847.
- [13] Felstein, Martin and Lawrence Summers 1979. "Inflation and the Taxation of Capital Income in the Corporate Sector." *National Tax Journal* 32(4), pp 445-470.
- [14] Flavin, Marjorie and Takashi Yamashita 2002. "Owner-Occupied Housing and the Composition of the Household Portfolio." *American Economic Review*, March, 345-62.
- [15] Geanakoplos, John, Michael Magill, and Martine Quinzii 2004. "Demography and the Long-run Predictability of the Stock Market." *Brookings Papers on Economic Activity* 1, pp. 241-307.
- [16] Grandmont, Jean-Michel 1977. "Temporary General Equilibrium." *Econometrica* 45 535-572.
- [17] 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.
- [18] 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.
- [19] Heaton, John and Deborah J. Lucas 2000. "Portfolio Choice and Asset Prices: The Importance of Entrepreneurial Risk." *Journal of Finance* 55, pp. 1163-1198
- [20] Hurst, Erik, Ming-Ching Luoh and Frank Stafford 1998. "Wealth Dynamics of American Families: 1984-1994. Brookings Papers on Economic Activity 1.
- [21] Krusell, Per and Anthony Smith 1998. "Income and Wealth Heterogeneity in the Macroeconomy." *Journal of Political Economy* 106(5). pp. 867-896.
- [22] McGrattan, Ellen R. and Edward C. Prescott 2005. "Taxes, Regulations, and the Value of U.S. and U.K. Corporations." Forthcoming *Review of Economic Studies*.
- [23] Piazzesi, Monika, Martin Schneider and Selale Tuzel 2005. "Housing, Consumption and Asset Pricing." Working paper, Chicago GSB, NYU & UCLA.
- [24] Summers, Lawrence H. 1981a. "Inflation, the Stock Market, and Owner-Occupied Housing." *American Economic Review* 71, 429-434.

- [25] Summers, Lawrence H. 1981b. "Inflation and the Valuation of Corporate Equities." *NBER working paper* 824.
- [26] Yao, Rui and Harold Zhang 2005. "Optimal Portfolio Choice with Risky Housing and Borrowing Constraints." *Review of Financial Studies* 18, pp. 197-239.