

Owner-Occupied Housing:

The Role of Collateral Constraints on the Household Portfolio

Marjorie Flavin
UCSD and NBER

Takashi Yamashita
H. Wayne Huizenga School of Business and Entrepreneurship
Nova Southeastern University

this version: October 2011

Corresponding author:
Marjorie Flavin
Department of Economics, 0508
University of California, San Diego
9500 Gilman Drive
La Jolla, CA 92093-0508
mflavin@ucsd.edu

ABSTRACT

The paper constructs a model of optimal portfolio allocation that focuses on the role of housing as collateral, allows for house price risk, and assumes that altering the quantity of housing incurs an adjustment cost. Because of the adjustment cost, the current house value becomes a state variable in the portfolio decision; the optimal portfolio depends on both the value of the housing state variable and the household's degree of relative risk aversion. Empirical results using the Survey of Consumer Finances support the model's implications regarding the effect of the housing state variable on the optimal portfolio of financial assets.

Until recently, the conventional wisdom in the portfolio choice literature held that the strategy of simply adding housing to the vector of assets and then constructing the optimal portfolio as the vector of asset holdings that achieves mean-variance efficiency was “incorrect.” As is often the case, the conventional wisdom is valid in a particular set of circumstances. In this case, the conventional wisdom -- that problems arise when applying a standard mean-variance optimization framework to a portfolio problem incorporating housing – is valid if the problem is set up as the choice over the quantities of all assets (the quantity of housing as well as the quantities of each financial asset) in a sequence of repeated mean-variance optimizations.

Housing cannot simply be treated symmetrically with financial assets in a portfolio problem because capital gains on housing are essentially different from capital gains on financial assets in the sense that an increase in the asset price of housing is strongly correlated with the price of a good which is quantitatively important in the household’s future consumption bundle (future housing services). House price appreciation, therefore, does not represent gains in wealth comparable to the gains that come from increases in the price of a financial asset.¹ For example, between 2000 and 2006, many U.S. homeowners experienced huge capital gains on their home (measured in dollar value). However, a \$200,000 increase in the value of a residence does not increase the household’s command over goods and services (i.e., is not an increase in wealth) in the same sense that a \$200,000 increase in the value of stock holdings would improve its command over goods and services, given that the household now faces a commensurately higher price for housing services.

The proscription against naively sticking housing into the vector of assets in a mean-variance framework holds when the quantity of housing is one of the choice variables in the optimization. In

¹ A forceful statement of the position that the change in the relative price of housing does not constitute a change in aggregate wealth is provided by Willem H. Buiter (2008), “Housing Wealth Isn’t Wealth.”

this paper, we consider the role of housing in a portfolio allocation problem by conditioning on the current holding of housing, and then finding the optimal holdings of financial assets.

Determining the optimal portfolio of financial assets conditional on housing, while treating the current level of housing as a state variable, is a well-defined subproblem within the household's overall optimization problem. That is, the overall problem of the household is to choose the optimal level of housing, holdings of financial assets, and the level of nondurable consumption in a continuous time setting. Adjustment of the quantity of housing requires the payment of a nonconvex adjustment cost, but nondurable consumption and financial assets can be adjusted frictionlessly. Because of the adjustment cost on housing, the solution to the general problem has a recursive structure: at each moment, the household considers whether or not to sell the house, pay the adjustment cost, and choose a new quantity of housing. Most of the time, it is not optimal to incur the adjustment cost. Having decided not to sell the house at that instant, the household then chooses the optimal level of nondurable consumption and the optimal holdings of financial assets conditional on the current level of housing. When, very infrequently, it is optimal to sell the house, the household optimally chooses the size of the new house. Thus, while the holding of housing is determined endogenously, it is not determined as part of the mean-variance optimization problem. Instead, the optimal holdings of financial assets, conditional on the current holdings of housing, is determined by the mean-variance framework.

The model is a variation of the housing model proposed in Flavin and Nakagawa (2008). Instead of assuming that the household can borrow or lend at the riskless rate and take negative as well as positive positions in all financial assets (as in Flavin and Nakagawa (2008)), in this paper we consider the portfolio allocation problem when the household is constrained by nonnegativity constraints on financial assets. In particular, we assume that the only way the household can borrow is to borrow against a house in the form of a mortgage, the size of the mortgage is limited to 100% of the value of the house, and that financial assets other than the mortgage can be held only in nonnegative

amounts. The constraint that the household can borrow only in the form of a mortgage is referred to as the borrowing, or collateral, constraint.

Incorporating the collateral and nonnegativity constraints considerably complicates the problem, and requires computational rather than analytic solution of the optimal portfolios. If it turned out that for most households for most of the time the constraints were not binding (i.e., households' optimal portfolios occurred at an interior solution despite the presence of the constraints), we could jettison the constrained version of the problem and work with the considerably simpler unconstrained version of the problem. To determine whether (and under what circumstances) the collateral and nonnegativity constraints are likely to be binding, we calculate the optimal portfolios for a range of assumptions on the stochastic structure of asset returns.

We then consider the implications of the model for the composition of the portfolio over the lifecycle. The model implies that, in the presence of the collateral and nonnegativity constraints, the optimal portfolio will depend on not only the household's degree of risk aversion, but also on the ratio of the house value to net worth. For a given degree of risk aversion, the percentage of the financial asset portfolio held in the form of stocks is a decreasing function of the ratio of house value to net worth over most of its range. Young homeowners typically have house values several times as large as their net worth; over the course of the lifecycle, the ratio of house value to net worth falls as the household accumulates wealth. Thus even if we consider two households with the same degree of risk aversion, the model predicts that the older household with a lower ratio of house value to net worth will generally hold a greater percentage of its portfolio of financial assets in the form of stocks than a younger household. Further, since the ratio of housing to net wealth varies across households, even within a given age group, the model introduces heterogeneity of financial portfolio holdings even within a group of households facing the same stochastic structure of the asset markets, with similar age, and with the identical preferences. In a similar spirit, Gomes and Michaelides (2005) incorporate

fixed costs and investor heterogeneity to model optimal portfolio allocation under the assumption that households have Epstein-Zin recursive utility and undiversifiable labor income risk. They establish that investor heterogeneity is the key to introducing life-cycle patterns of risky-asset holdings observed in the data. Our model, on the other hand, implies that the size of housing investment relative to net worth is an important determinant of an investor's optimal portfolios.

Finally, we conclude by presenting summary statistics, graphs, and regression results documenting the lifecycle patterns in household portfolios using the repeated cross sections provided by the Survey of Consumer Finances (SCF). The analytical model implies that the household's optimal portfolio will depend on the household's degree of risk aversion and on the state variable representing the ratio of house value to net worth. The numerical results indicate that, for a given degree of risk aversion, the optimal portfolios will systematically vary with the ratio of house value to net worth. In addition to providing the quantitative data on asset holding, the SCF contains several questions designed to elicit the respondents' tolerance for risk. Based on the responses to these questions on risk tolerance, households can be categorized as "risk tolerant", "average risk aversion", or "risk averse", and these (admittedly rough) measures of risk preferences are used to control for risk aversion in the statistical tests of the model.

There are many channels through which housing can alter a household's risk-return trade off and hence influence the optimal financial portfolio (see, for example, Grossman and Laroque (1990), Flavin and Yamashita (2002), Cocco (2005), Hu (2005), Yao and Zhang (2005), and Cauley, Pavlov, and Schwartz (2007)). However, many studies adopt simplifying assumptions that may not reflect the characteristics of housing investment. For example, Cocco (2005) assumes that the value of the home is perfectly correlated with aggregate labor income shocks, thus is non-stochastic with respect to permanent income. Yao and Zhang (2005), while endogenizing housing tenure decisions, assume unit elasticity between housing and non-durable consumption, which is much higher than the available

estimates. They calibrate the optimal portfolio holdings using the risk premium on risky assets of 4 percent and the zero expected return on housing investment. Thus stocks are less attractive than the historical average and a household would purchase a home purely for consumption purposes.

Also related is the Faig and Shum (2002) model of personal illiquid projects. In their model, illiquidity in a personal project requires the investor to maintain relatively safe and liquid financial portfolios. The two models are similar as the assumed lack of correlation between personal projects and financial asset returns rules out the possibility of hedging the investment in illiquid assets with financial portfolios. However, their analogy of housing as an example of illiquid personal projects does not conform to the characteristics of the housing markets. Faig and Shum assume that investors would incur substantial losses when a house is sold before the final period. However, given that secondary markets for housing are extremely well developed, the personal illiquid projects in their model should probably be interpreted as small businesses rather than housing. Our model is specifically tailored to investigate the role of housing in asset allocation, as we model explicitly housing price risk, the role of housing services as a consumption good, and the infrequent adjustment of housing. Further, instead of treating housing primarily as an additional source of background risk, our model focuses on the role of housing in providing collateral.²

Section I: A model of asset allocation with housing as an asset

In analyzing the role of housing in the portfolio allocation problem, our objective is to model housing in a way that incorporates the important aspects in which housing differs from financial assets. To that end, we assume that, because of imperfections in the rental market for housing services, the homeowner's decision regarding the quantity of residential real estate to acquire simultaneously determines both its investment in housing as an asset and its consumption of housing services. As in

² Small business entrepreneurs often use housing as collateral to finance their businesses. See Hurst and Lusardi (2004).

Grossman and Laroque (1990) and Flavin and Nakagawa (2008), once the household purchases a particular house, no adjustments to its size (or any other attribute) can be made without selling the existing house, incurring an adjustment cost proportional to the value of the house sold, and purchasing a new house. Because the transactions cost is assumed proportional (with factor of proportionality λ) to the value of the house sold, the model incorporates a “nonconvex” adjustment cost.

The instantaneous utility function depends (nonseparably) on housing services, which is assumed proportional to the stock of housing, and on a second good, referred to as the “nondurable consumption good,” denoted C_t , which is costlessly adjustable. The household’s expected lifetime utility is given by:

$$(1) \quad U = E_t \int_0^{\infty} e^{-\delta t} u(H_t, C_t) dt$$

The notation H_t represents a physical measure of the quantity of housing; in the simplest specification H_t can be thought of as a scalar measure of the square footage of the home, but in a more elaborate specification H_t could be interpreted as a vector of physical characteristics (square footage, number of fireplaces, quality of finish materials, etc). The crucial point is that H_t reflects some physical measure of the quantity of housing rather than the market value of the house. For ease of exposition, we interpret H_t as a scalar reflecting the square footage of the house. Using the nondurable good as numeraire, define:

P_t = house price (per square foot) in the household's current market

(2) P'_t = house price (per square foot) in the region to which the household relocates in the next move

Housing is subject to capital gains and losses, in the sense that the price of housing relative to the second consumption good is assumed to vary over time. Further, the model allows for cross sectional variation in the price of housing; housing prices in two regional markets will presumably be

correlated but are not necessarily perfectly correlated.

In the Grossman and Laroque model, only “endogenous” moves are considered as the household sells the current house and purchases a different one only when the endogenously determined state variable, the ratio of wealth to housing, hits the lower or upper S,s bounds. Their numerical results indicate that the mean time between moves is on the order of 25 to 40 years, depending on the choice of parameter values. Given that households move much more frequently than can be explained by endogenous moves alone, we assume that some, perhaps most, moves are precipitated by factors exogenous to the model, such as job relocation, or changes in household composition due to marriage or divorce. Taken literally, the model assumes that when the current house is sold, the household is exogenously assigned to a different regional real estate market. While allowing the household to choose the regional real estate market of its new home purchase would add to the realism of the model, endogenizing the choice the subsequent regional market is beyond the scope of the paper.

While the assumption that, upon sale of the current house, the household is exogenously assigned to the subsequent regional market is invoked primarily for tractability, we argue that the exogeneity assumption is plausible as an interpretation of the bequest motive. If we think of the infinite horizon household as a family dynasty, the death of one generation constitutes an exogenous event which precipitates the sale of the current house. The proceeds of the house sale are distributed to the next generation, who are distributed exogenously (from the point of view of the generation making the bequest) across the country. In this dynastic interpretation of the household, bequests are stochastic (because death is stochastic), and the generation currently alive values the size of the bequest in terms of its purchasing power for the recipients, but takes as exogenous the regional location of the recipients.

In addition to housing, the household can invest in any of n risky financial assets, including T-

bills, bonds, and stocks. There is no riskless asset, although the risk involved in holding T-bills is very small. Households can borrow, in the form of mortgages, an amount up to the value of their home. Unlike housing, financial assets (including mortgages) can be bought and sold with zero transaction cost. We abstract from labor income or human wealth, and assume that wealth is held only in the form of financial assets and housing. Wealth is thus given by:

$$(3) \quad W_t = P_t H_t + \underline{X}_t \underline{\ell}$$

where $\underline{X}_t = (1 \times n)$ vector of amounts (expressed in terms of the nondurable good) of the risky assets held, and $\underline{\ell} = (n \times 1)$ vector of ones. Using the first element of $\underline{X}_t = (1 \times n)$ to represent the mortgage, the corner constraints on the vector of financial assets are given by:

$$(4a) \quad \begin{aligned} -P_t H_t \leq X_{1t} \leq 0 \\ X_{1t} \leq W_t - P_t H_t \end{aligned} \quad (\text{collateral constraint on mortgage borrowing})$$

$$(4b) \quad 0 \leq X_{it} \quad i = 2 \text{ to } n \quad (\text{nonnegativity constraints on other financial assets})$$

Equations (4a) and (4b) reflect our assumption that the household can borrow against the house but not against financial assets or sell financial assets short. By imposing the collateral constraint on mortgage borrowing (4a), and the nonnegativity constraints on other financial assets (4b), we depart from the housing model of Flavin and Nakagawa (2008), which assumed that households could borrow or lend at the riskless rate. In this paper, because we are interested in characterizing the portfolio behavior of the typical, or median household, we impose the collateral constraint and nonnegativity constraints in order to model the realistic market constraints faced by a typical household.

Assuming that interest and dividend payments are reinvested so that the total return is received in the form of appreciation of the value of the asset, let $b_{i,t}$ denote the value of the i th risky asset. The vector of prices of the risky financial assets follows an n -dimensional Brownian motion process:

$$(5) \quad db_{i,t} = b_{i,t} ((\mu_i + r_f)dt + d\omega_{i,t})$$

Define the vector $\underline{\omega}_{F,t} = (\omega_{1,t}, \omega_{2,t}, \dots, \omega_{n,t})$ as an n-dimensional Brownian motion with zero drift and with instantaneous covariance matrix Σ . Also define the corresponding vector of expected returns on financial assets as $\underline{\mu} = (\mu_1, \mu_2, \dots, \mu_n)$. House prices also follow a Brownian motion:

$$(6) \quad \begin{aligned} dP_t &= P_t (\mu_H dt + d\omega_{Ht}) \\ dP'_t &= P'_t (\mu_H dt + d\omega_{H't}) \end{aligned}$$

where ω_{Ht} and $\omega_{H't}$ are Brownian motions with zero drift, instantaneous variance σ_P^2 and $\sigma_{P'}^2$, respectively, and instantaneous covariance σ_H .

Combining equations (5) and (6), define the $((n+2) \times 1)$ vector

$$(7) \quad d\underline{\omega}_t = \begin{bmatrix} d\omega_{1t} \\ \square \\ d\omega_{nt} \\ d\omega_{Ht} \\ d\omega_{H't} \end{bmatrix}$$

which has instantaneous $((n+2) \times (n+2))$ covariance matrix Ω :

$$(8) \quad \Omega = \begin{bmatrix} \Sigma & 0 & 0 \\ 0 & \sigma_P^2 & \sigma_H \\ 0 & \sigma_H & \sigma_{P'}^2 \end{bmatrix}$$

By specifying Ω as a block diagonal matrix, the model imposes the assumption that the stochastic component of house prices, both in the current market and in the household's next market, is uncorrelated with the returns to any of the financial assets. Note that no restrictions are imposed on σ_H , the covariance of house prices in the current market with house prices in the household's next market. While the analytical results concerning the composition of the optimal portfolio require the block diagonality of the covariance matrix Ω , this assumption does not impose the (extremely implausible) assumption that house prices movements are uncorrelated across regions.

To characterize the household's maximization problem, let $V(H, W, P, P')$ denote the

supremum of household expected utility, conditional on initial conditions (H, W, P, P') .

$$(9) \quad V(H_0, W_0, P_0, P'_0) = \sup_{\underline{X}_s, C_s, \tau_1} E \left[\int_0^{\tau_1} e^{-\delta s} u(H_0, C_s) ds + e^{-\delta \tau_1} V(H_{\tau_1}, W_{\tau_1}, P_{\tau_1}, P'_{\tau_1}) \right]$$

At any moment, the household decides whether to “stop”, i.e., incur the transactions cost and sell the current house. Optimal stopping times are denoted $\tau_1, \tau_2, \tau_3, \dots$. At any stopping time, the household chooses the size of the new house in order to maximize expected utility. Between stopping times, when the level of housing is fixed, the household chooses the path of nondurable consumption and the path of financial asset holdings. We are primarily interested in the household’s determination of nondurable consumption and financial asset holdings during a short time interval $(0, t)$ within which stopping does not occur. During such a time interval, wealth evolves according to:

$$(10) \quad dW_t = \left[P_t H_0 \mu_H + \underline{X}_t \underline{\mu} - C_t \right] dt + \underline{X}_t d\omega_{Ft} + P_t H_0 d\omega_{Ht}$$

and the Bellman equation is:

$$(11) \quad V(H_0, W_0, P_0, P'_0) = \sup_{\underline{X}_s, C_s} E \left[\int_0^t e^{-\delta s} u(H_0, C_s) ds + e^{-\delta t} V(H_0, W_t, P_t, P'_t) \right]$$

subject to the budget constraint (10) and the process for house prices (6). Subtracting

$V(H_0, W_0, P_0, P'_0)$ from both sides, dividing by t and taking the limit as $t \rightarrow 0$ gives:

$$(12) \quad 0 = \lim_{t \rightarrow 0} \sup_{\underline{X}_s, C_s} E \left[\frac{1}{t} \int_0^t e^{-\delta s} u(H_0, C_s) ds + \frac{1}{t} \left(e^{-\delta t} V(H_0, W_t, P_t, P'_t) - V(H_0, W_0, P_0, P'_0) \right) \right]$$

Evaluating the integral and using Ito’s lemma, equation (12) can be rewritten as:

$$(13) \quad 0 = \sup_{\underline{X}_0, C_0} \left\{ u(H_0, C_0) - \delta V(H_0, W_0, P_0, P'_0) + \frac{\partial V}{\partial W} (P_0 H_0 \mu_H + \underline{X}_0 \underline{\mu} - C_0) + \frac{\partial V}{\partial P} P_0 \mu_H \right. \\ \left. + \frac{\partial V}{\partial P'} P_0 \mu_{H'} + \frac{1}{2} \frac{\partial^2 V}{\partial W^2} (\underline{X}_0 \Sigma \underline{X}_0^T + P_0^2 H_0^2 \sigma_P^2) + \frac{1}{2} \frac{\partial^2 V}{\partial P^2} P_0^2 \sigma_P^2 + \frac{1}{2} \frac{\partial^2 V}{\partial P'^2} P_0'^2 \sigma_{P'}^2 \right. \\ \left. + \frac{\partial^2 V}{\partial W \partial P} P_0^2 H_0 \sigma_P^2 + \frac{\partial^2 V}{\partial W \partial P'} P_0 P_0' H_0 \sigma_H + \frac{\partial^2 V}{\partial P \partial P'} P_0 P_0' \sigma_H \right\}$$

Because nondurable consumption is assumed to be costlessly adjustable, the household equates the

marginal utility of nondurable consumption with the marginal value of wealth:

$$(14) \quad \frac{\partial u}{\partial C} = \frac{\partial V}{\partial W}$$

Only two of the terms in equation (13) actually depend on financial asset holdings, \underline{X}_0 . Thus the household chooses its portfolio of financial assets according to the rule:

$$(15) \quad \sup_{\underline{x}_0} \left\{ \frac{\partial V}{\partial W} (P_0 H_0 \mu_H + \underline{X}_0 \underline{\mu} - C_0) + \frac{1}{2} \frac{\partial^2 V}{\partial W^2} (\underline{X}_0 \Sigma \underline{X}_0^T + P_0^2 H_0^2 \sigma_P^2) \right\}$$

Restating financial asset holdings and the value of the house as shares of current wealth,

$$(16) \quad \begin{aligned} \underline{x} &= \frac{\underline{X}_0}{W_0} \\ h &= \frac{P_0 H_0}{W_0} . \end{aligned}$$

The optimization problem can be rewritten, after including the term $\frac{\partial V}{\partial W} C_0$ in the constant term, as

$$(17) \quad \sup_{\underline{x}} \left\{ \frac{\partial V}{\partial W} W_0 (h \mu_H + \underline{x} \underline{\mu}) + \frac{1}{2} \frac{\partial^2 V}{\partial W^2} W_0^2 (\underline{x} \Sigma \underline{x}^T + h^2 \sigma_P^2) \right\}$$

Thus the household chooses asset shares, \underline{x} , in order to maximize:

$$(18) \quad \text{objective function} = (h \mu_H + \underline{x} \underline{\mu}) - \frac{1}{2} A_0 (\underline{x} \Sigma \underline{x}^T + h^2 \sigma_P^2)$$

subject to the constraint

$$(19) \quad 1 = h + \underline{x} \underline{1}$$

and the nonnegativity constraints (equation (4b)) on the elements of \underline{x} . In equation (18), A_0 denotes

the household's degree of relative risk aversion:

$$A_0 \equiv - \frac{\frac{\partial^2 V(H_0, W_0, P_0, P_0')}{\partial W_0^2} W_0}{\frac{\partial V(H_0, W_0, P_0, P_0')}{\partial W_0}} \geq 0$$

Equation (18) states that the household's objective function is an increasing function of the expected return, $\underline{x}\underline{\mu}$, and a decreasing function of the variance, $\underline{x}\underline{\Sigma}\underline{x}^T$, of the portfolio of financial assets. Thus we can interpret equation (18) as saying that the optimal choice of \underline{x} will be on the mean-variance efficient frontier of financial assets. The implication that the optimal portfolio will be mean-variance efficient does not require a specific assumption such as constant relative risk aversion on the instantaneous utility function.

The derivation of equation (18) required that the covariance matrix be block diagonal as specified in equation (8). The dependence of the mean-variance efficiency result on the assumption of block-diagonality can be understood intuitively. Due to the transactions costs associated with selling the house, the optimization problem has the following recursive structure. The household first considers whether it is optimal to sell the house immediately, i.e., considers whether $t = 0$ is a stopping time. If $t = 0$ is not a stopping time, the household has decided to hold $P_t H_t$ in the form of housing for (at least) this instant and is therefore subject to an instantaneous expected return and standard deviation of return on the house as determined by the parameters μ_H and σ_P^2 . If the covariance matrix is block diagonal, returns to financial assets are uncorrelated with current or future house prices. In this case, even though the risk averse household will dislike the risk created by variability in current (P) or future (P') house prices, the household is unable to hedge either of these types of risk with the portfolio of financial assets. Since financial assets cannot be used to hedge the risks associated with changes in current or future house prices, the model implies that the optimal vector of financial assets will achieve mean variance efficiency with respect to the portfolio of financial assets.

Although the model has the implication that the optimal vector of financial assets holdings will be mean-variance efficient, this result does not imply that the household's optimal portfolio is independent of the current holding of housing assets, or of the level of housing prices. The state

variables that characterize the housing sector (H_t , P_t , and P_t') influence the optimal portfolio of financial assets in two ways: first, in determining the location of the constrained mean-variance efficient frontier available to the household, and second, in determining the household's degree of risk aversion (A_t) and thus its optimal location on the constrained frontier. The effect of the housing state variables on the location of the constrained efficient frontier is a result of the assumptions that households can borrow only in the form of a mortgage and the size of the mortgage is limited to 100% of the value of the house. That is, the household optimizes its portfolio of financial assets subject to the collateral constraint given in equation (4a): a household whose house value exceeds net worth must hold a mortgage with minimum size equal to $[P_t H_t - W_t]$ and with maximum size equal to $[P_t H_t]$. Because the minimum and maximum constraints on the holdings of one of the financial assets (the mortgage) depend directly on the house value, $P_t H_t$, the constrained mean-variance efficient frontier available to the household also depends on the house value.

If, as in Flavin and Nakagawa (2008), we dropped the collateral and nonnegativity constraints and simply assumed interior solutions for every element of \underline{x} , we could differentiate equation (18) with respect to \underline{x} and obtain an analytical solution for the portfolio shares

$$(20) \quad \underline{x} = \begin{bmatrix} -\frac{\partial V}{\partial W} \\ \frac{\partial^2 V}{\partial W^2} W_0 \end{bmatrix} \Sigma^{-1} \underline{\mu}$$

which implies that all households hold risky assets in the same proportions, the mutual fund separation theorem holds, and that the CAPM holds. However, under the current assumption that the household faces the constraints in equation (4), the possibility of corner solutions for some of the elements of \underline{x} implies that equation (18) requires numerical optimization.

The two terms in parentheses on the right hand side of equation (18) represent the expected

return and the variance of the asset portfolio inclusive of housing. Because the covariance matrix is block diagonal and h is a state variable, any vector \underline{x} that achieves mean-variance efficiency with respect to the portfolio of financial assets also achieves mean-variance efficiency with respect to the whole portfolio inclusive of housing. Thus the model not only implies that the optimal portfolio of financial assets will be mean-variance efficient; it further implies that the mean-variance efficiency property applies both to the portfolio of financial assets and to the portfolio inclusive of housing. Whether we choose to think about the efficient frontier in terms of the expected return and standard deviation of the portfolio inclusive of housing or in terms of the expected return and standard deviation of the portfolio of financial assets alone, the state variable h affects the efficient frontier via the corner constraints.

Housing also affects a household's choice of optimal portfolio through its effect on the curvature of the value function. From among the set of optimal portfolios on the constrained efficient frontier, the optimal portfolio is determined by the household's tradeoff between risk and return as represented by the curvature of the value function, A_t .

$$(21) \quad A_t \equiv - \frac{\frac{\partial^2 V(W_t, H_t, P_t, P_t')}{\partial W_t^2}}{\frac{\partial V(W_t, H_t, P_t, P_t')}{\partial W_t}} W_t > 0$$

In general, the curvature of the value function will depend on the values of all of the state variables.⁴

Thus the optimization problem of the household can be written as:

$$(22) \quad \sup_{\underline{x}} \left(\mu - \frac{A_t}{2} \sigma^2 \right) \quad \text{subject to the constraints (4) and (19) and}$$

$$(23) \quad \mu \equiv h\mu_H + \underline{x}\underline{\mu} \equiv \text{expected return on portfolio inclusive of housing, and}$$

⁴ In a related paper, Flavin (2009) obtains the household's degree of relative risk aversion as a function of the state for a simplified version of the model by (numerically) solving the household's optimal stopping problem.

$$(24) \quad \sigma^2 \equiv \underline{x}\Sigma\underline{x}^T + h^2\sigma_p^2 \equiv \text{variance of return on portfolio inclusive of housing.}$$

From equation (22), the slope of the household's indifference curve is:

$$(25) \quad \frac{\partial \mu}{\partial \sigma} = A_t \sigma$$

With the constrained efficient frontier and the indifference curve, we can identify the household's optimal portfolio as a function of its constraints, as measured by h , and its degree of risk aversion, as measured by the curvature of the value function, A_t .

Section II: Optimal portfolios as a function of risk aversion and the housing constraint

For tractability, the housing model in Section I relies on the assumption that the covariance matrix is block diagonal. In a previous paper, Flavin and Yamashita (2002), we estimated the mean return, and the covariance matrix of returns, to housing, a mortgage, T-bills, T-bonds, and stocks, using household level data from the PSID from 1968-1992. Table I reports the expected returns, covariance matrix, and correlation matrix based on the PSID data.

Table I: Expected Returns and Covariance Matrix – PSID Data

	T-Bills	Bonds	Stocks	Mortgage	House
Mean Return (arithmetic)	-.0038	.0060	.0824	.0000	.0659
Standard Deviation	.0435	.0840	.2415	.0336	.1424
Covariance Matrix					
T-Bills	.0018920				
T-Bonds	.0025050	.0070613			
Stocks	.0002008	.0040381	.0583292		
Mortgage	.0007087	.0023854	.0025400	.0011274	
House	-.000119	-.000067	-.000178	-.0000057	.020284
Correlation Matrix					
T- Bills	1.0000				
T-Bonds	.68533 (.09103)	1.0000			
Stocks	.01912 (.12498)	.19897 (.12251)	1.0000		
Mortgage	.84119 (.11529)	.680286 (.15626)	.467954 (.18842)	1.0000	
House	-.03339	-.004506	-.000771	-.001192	1.0000

(.21309) (.21320) (.21319) (.21320)

Source: Flavin and Yamashita (2002). Standard errors are in parentheses.

According to the bottom row of the correlation matrix, the assumption that the covariance matrix is block diagonal in the sense that the return to housing is uncorrelated with the return to each of the financial assets is fully consistent with the data from the PSID. In each case, the correlation of the return to T-bills, T-bonds, stocks, and mortgages has a correlation with the return to housing which is essentially zero both in terms of numerical size and statistical significance.

While the historical data is valuable for testing the validity of the block-diagonality assumption, examination of the vector of mean returns over this sample period provides an illustration of the distinction between sample moments and population moments. Ex post, the average after-tax, real return on T-bills was slightly negative, the average after-tax rate on mortgages was zero (to four decimal places!), and the return to Treasury bonds was only 60 basis points. While these statistics accurately characterize the historical returns, ex post, it seems unlikely that actual households were making their portfolio decisions based on the ex ante belief that the average returns to these nominal assets would be so low.

Instead of using the sample moments, we attempt to write down the subjective assumptions on the risk and return on which we base on own household portfolio decisions. Further, we calculate the optimal portfolios for several different sets of assumptions on the moments of asset returns to check the robustness of the results. By varying the assumptions on the moments of asset returns, we can cover most of the specifications commonly used in the literature.

The baseline set of assumptions is reported in Table IIa; the after-tax real return on T-bills is assumed to be small but positive (0.01), the returns on bonds and mortgages are equal at 0.03, and the return on stocks (0.07) is slightly higher than the return on housing (0.05). For the baseline case, the assumed covariance matrix of returns is a rough approximation to the covariance matrix estimated from the PSID, although the numerical values are limited to two significant digits.

The optimal portfolios calculated under the baseline assumptions are reported in Table IIb. The table reports the optimal holdings of T-bills, T-bonds, and stocks as percentages of the portfolio of financial assets and the size of the mortgage is expressed as a percent of the house value (a mortgage value of -1 reflects a 100% mortgage). Thus any cell that reports a value of unity or zero for the share of financial assets, or a negative one for the mortgage, represents a portfolio in which at least one of the corner constraints is binding.

The nonnegativity constraint on T-bills is almost always binding. Only when total net worth is twice the value of the house and risk aversion is high, does the optimal portfolio contain a strictly positive amount of T-bills. For households that are highly risk tolerant (with relative risk aversion of unity), the optimal strategy is to borrow the maximum against the house, and put all of the household's net worth into stocks, independent of the value of h . For a given value of h , higher values of relative risk aversion induce the household to decrease the share of the portfolio held in stocks, and at the same time reduce leverage by reducing the loan-to-value ratio on the house.

For a given value of risk aversion, as the value of the h declines the optimal portfolio is characterized by a lower loan-to-value ratio, and, in general, an increase in the share of the portfolio devoted to stocks. The dependence of the portfolio share devoted to stocks is not monotonic in h over the whole range, however. If the ratio of house value to net worth is less than one, the optimal share devoted to stocks declines with further increases in h for moderate and high levels of risk aversion.

Table IIIa states an alternative set of assumptions on the stochastic process of asset returns. Here stocks are assumed to have a higher expected return (0.09 instead of 0.07) and higher standard deviation of return (0.25 instead of 0.20), while housing is assumed to have a lower expected return (0.03 instead of 0.05) and lower standard deviation (0.10 instead of 0.15). The resulting optimal portfolios are reported in Table IIIb. Comparison of Tables 2b and 3b indicates that the quantitative effect on the optimal portfolio shares is modest.

A third set of assumption is considered in Table IVa. Here the expected return and standard deviation of stocks is lower than in the baseline case (expected return reduced from 0.07 to 0.05 and standard deviation of return reduced from 0.20 to 0.15), while the expected return and standard deviation of returns to housing are increased (mean return increased from 0.05 to 0.07 and standard deviation increased from 0.15 to 0.20). Again, the optimal portfolios generated under the new set of assumptions are not dramatically different from those generated from the baseline assumptions. The portfolios generated by any of the three sets of assumptions conform to the same set of qualitative characteristics: First, the nonnegativity constraint on T-bills is almost always binding, second, the share of the portfolio held in the form of stocks is decreasing in the value of h over most of its range.

Table IIa: Baseline assumptions on mean returns and covariance matrix of returns

	T-Bills	Bonds	Stocks	Mortgage	House
Mean Return (arithmetic)	0.01	0.03	0.07	0.03	0.05
Standard Deviation	0.04	0.10	0.20	0.04	0.15
Covariance Matrix					
T-Bills	0.0016				
T-Bonds	0.0025	0.010			
Stocks	0.0005	0.005	0.040		
Mortgage	0.0010	0.003	0.003	0.0016	
House	0	0	0	0	0.0225

Table IIb: Optimal Portfolio Weights for Different Constraints on h

Housing-to NW Ratio	Assets in Portfolio	Curvature of value function, A				
		$A = 1$	$A = 2$	$A = 4$	$A = 8$	$A = 10$
3.50	Treasury Bills	0	0	0	0	0
	Treasury Bonds	0	0.3750	0.6132	0.7242	0.7506
	Stocks	1	0.6250	0.3868	0.2758	0.2494
	Mortgage	-1	-1	-0.9871	-0.9512	-0.9440
3.00	Treasury Bills	0	0	0	0	0
	Treasury Bonds	0	0.3750	0.5679	0.6934	0.7242
	Stocks	1	0.6250	0.4321	0.3066	0.2758
	Mortgage	-1	-1	-0.9380	-0.8961	-0.8878
2.50	Treasury Bills	0	0	0	0	0
	Treasury Bonds	0	0.3342	0.5037	0.6468	0.6837
	Stocks	1	0.6658	0.4927	0.3532	0.3163
	Mortgage	-1	-0.9698	-0.8693	-0.8191	-0.8090
2.00	Treasury Bills	0	0	0	0	0
	Treasury Bonds	0	0.2372	0.4057	0.5679	0.6132
	Stocks	1	0.7628	0.5943	0.4321	0.3868
	Mortgage	-1	-0.8920	-0.7663	-0.7035	-0.6910
1.50	Treasury Bills	0	0	0	0	0
	Treasury Bonds	0	0.0977	0.2372	0.4057	0.4640
	Stocks	1	0.9023	0.7628	0.5943	0.5396
	Mortgage	-1	-0.7621	-0.5946	-0.5109	-0.4941
1.00	Treasury Bills	0	0	0	0	0
	Treasury Bonds	0	0	0	0	0
	Stocks	1	1	1	1	1
	Mortgage	-1	-0.5618	-0.2809	-0.1404	-0.1124
0.75	Treasury Bills	0	0	0	0	0
	Treasury Bonds	0	0	0	0.3750	0.4750
	Stocks	1	1	1	0.6250	0.5250
	Mortgage	-1	-0.4026	-0.0281	0	0
0.50	Treasury Bills	0	0	0	0.2157	0.3961
	Treasury Bonds	0	0	0.3750	0.4255	0.3086
	Stocks	1	1	0.6250	0.3588	0.2953
	Mortgage	-1	-0.0843	-0.0281	0	0

Note: Shares of T-bills, bonds, and stocks are stated as a percentage of the portfolio of financial assets, so that for each portfolio the shares of these three assets must sum to one. The mortgage is expressed as a percent of the house value, i.e., Mortgage = -1 indicates a 100% mortgage.

Table IIIa: Relative to baseline, higher mean and s.d. of stocks; lower mean and s.d. of house

	T-Bills	Bonds	Stocks	Mortgage	House
Mean Return (arithmetic)	0.01	0.03	.07 to 0.09	0.03	.05 to 0.03
Standard Deviation	0.04	0.10	.20 to 0.25	0.04	.15 to 0.10
Covariance Matrix					
T-Bills	0.0016				
T-Bonds	0.0025	0.010			
Stocks	0.0005	0.005	0.0625		
Mortgage	0.0010	0.003	0.003	0.0016	
House	0	0	0	0	0.010

Table IIIb: Optimal Portfolio Weights for Different Constraints on h

Housing-to-NW Ratio	Assets in Portfolio	Curvature of value function, A				
		$A = 1$	$A = 2$	$A = 4$	$A = 8$	$A = 10$
3.50	Treasury Bills	0	0	0	0	0
	Treasury Bonds	0	0.4400	0.6545	0.7678	0.7946
	Stocks	1	0.5600	0.3455	0.2322	0.2054
	Mortgage	-1	-1	-0.9725	-0.9396	-0.9330
3.00	Treasury Bills	0	0	0	0	0
	Treasury Bonds	0	0.4400	0.6080	0.6934	0.7678
	Stocks	1	0.5600	0.3920	0.2635	0.2322
	Mortgage	-1	-1	-0.9231	-0.8846	-0.8769
2.50	Treasury Bills	0	0	0	0	0
	Treasury Bonds	0	0.3653	0.5418	0.6889	0.7265
	Stocks	1	0.6353	0.4582	0.3111	0.2735
	Mortgage	-1	-0.9462	-0.8539	-0.8077	-0.7985
2.00	Treasury Bills	0	0	0	0	0
	Treasury Bonds	0	0.2632	0.4400	0.6080	0.6545
	Stocks	1	0.7368	0.5600	0.3920	0.3455
	Mortgage	-1	-0.8654	-0.7500	-0.6923	-0.6808
1.50	Treasury Bills	0	0	0	0	0
	Treasury Bonds	0	0.1148	0.2632	0.4400	0.4694
	Stocks	1	0.8852	0.7368	0.5600	0.5031
	Mortgage	-1	-0.7308	-0.5769	-0.500	-0.4941
1.00	Treasury Bills	0	0	0	0	0
	Treasury Bonds	0	0	0	0	0
	Stocks	1	1	1	1	1
	Mortgage	-1	-0.5164	-0.2582	-0.1291	-0.1033
0.75	Treasury Bills	0	0	0	0	0
	Treasury Bonds	0	0	0	0.4400	0.5360
	Stocks	1	1	1	0.5600	0.4640
	Mortgage	-1	-0.3471	-0.0029	0	0
0.50	Treasury Bills	0	0	0	0.2157	0.3961
	Treasury Bonds	0	0	0.4400	0.4529	0.3322
	Stocks	1	1	0.5600	0.3086	0.2520
	Mortgage	-1	-0.0086	0	0	0

Table IVa: Relative to baseline, lower mean and s.d. of stocks; higher mean and s.d. of house

	T-Bills	Bonds	Stocks	Mortgage	House
Mean Return (arithmetic)	0.01	0.03	.07 to 0.05	0.03	.05 to 0.07
Standard Deviation	0.04	0.10	.20 to 0.15	0.04	.15 to 0.20
Covariance Matrix					
T-Bills	0.0016				
T-Bonds	0.0025	0.010			
Stocks	0.0005	0.005	0.0225		
Mortgage	0.0010	0.003	0.003	0.0016	
House	0	0	0	0	0.040

Table IVb: Optimal Portfolio Weights for Different Constraints on h

Housing-to-NW Ratio	Assets in Portfolio	Curvature of value function, A				
		$A = 1$	$A = 2$	$A = 4$	$A = 8$	$A = 10$
3.50	Treasury Bills	0	0	0	0	0
	Treasury Bonds	0	0.3333	0.5556	0.6548	0.6766
	Stocks	1	0.6667	0.4444	0.3452	0.3243
	Mortgage	-1	-1	-1	-0.9724	-0.9653
3.00	Treasury Bills	0	0	0	0	0
	Treasury Bonds	0	0.3333	0.5227	0.6291	0.6548
	Stocks	1	0.6667	0.4773	0.3709	0.3452
	Mortgage	-1	-1	-0.9571	-0.9158	-0.9076
2.50	Treasury Bills	0	0	0	0	0
	Treasury Bonds	0	0.3162	0.4671	0.5900	0.6210
	Stocks	1	0.6838	0.5329	0.4100	0.3790
	Mortgage	-1	-0.9851	-0.8861	-0.8366	-0.8267
2.00	Treasury Bills	0	0	0	0	0
	Treasury Bonds	0	0.22270	0.3805	0.5227	0.5614
	Stocks	1	0.7730	0.6195	0.4773	0.4386
	Mortgage	-1	-0.9035	-0.7797	-0.7178	-0.7054
1.50	Treasury Bills	0	0	0	0	0
	Treasury Bonds	0	0.0951	0.2270	0.3805	0.4291
	Stocks	1	0.9049	0.7730	0.6195	0.5709
	Mortgage	-1	-0.7621	-0.6023	-0.5198	-0.5033
1.00	Treasury Bills	0	0	0	0	0
	Treasury Bonds	0	0	0	0	0
	Stocks	1	1	1	1	1
	Mortgage	-1	-0.5525	-0.2762	-0.1381	-0.1105
0.75	Treasury Bills	0	0	0	0	0
	Treasury Bonds	0	0	0	0.3333	0.4222
	Stocks	1	1	1	0.6667	0.5778
	Mortgage	-1	-0.3775	-0.0092	0	0
0.50	Treasury Bills	0	0	0	0.1882	0.3710
	Treasury Bonds	0	0	0.3333	0.3925	0.2785
	Stocks	1	1	0.6667	0.4193	0.3505
	Mortgage	-1	-0.0276	0	0	0

Section III: Cross-Year, Cross-Section, and Cohort Analysis

Because net wealth typically rises more dramatically than house value over the household's working years, for most households the ratio of house value to net worth declines over the life-cycle. Since our model predicts that the share of the portfolio held in stocks generally increases as the ratio of house value to net worth (h) declines, the dependence of the optimal portfolio on h will induce a life-cycle pattern in portfolio composition. In this section we examine the empirical relationships among the house-value-to-net-worth ratio (h), the stock-to-financial-assets ratio (s), and the loan-to-value (LTV) ratio using the seven waves of the Survey of Consumer Finances (SCF). In this analysis, we examine the relationship between the variables of interest both in the cross section, and in synthetic cohorts constructed from the repeated cross sections of the SCF. The sample period includes two recessions (1990-1991 and 2001) and two asset-price booms, the dot-com boom in the stock markets led by the technology stocks (1995-2001) and the housing boom (2001-2006). Household balance sheets are constructed using the program provided by the Board of Governors in its SCF web page (Federal Reserve Board 2008).⁵ The sample is limited to households with heads between age of 24 and 89 at the time of the survey.⁶

Table V summarizes the key sample statistics of the evolution of American households' wealth holdings between 1989 and 2007. All dollar amounts are expressed in constant 2007 dollars. Median household net worth increased from \$100,800 in 1989 to \$154,000 in 2007, with a decline after the 1991 recession and a peak in 2001 after the stock market boom. The

⁵ The SCF imputes missing values of responses to the survey questions. We use all five imputations and take arithmetic means of the five replicates. We then use the main replicate weight to arrive at estimates of sample statistics such as means and medians.

⁶ We also limit our sample to observations that have non-negative net worth.

median value of the principal residence, which rose gradually from \$120,700 in 1989 to \$128,600 in 1998, increased rapidly after 1998 to reach \$200,400 in 2007.

The quantitative importance of housing is reflected in the change in the value of h (house value as a percentage of net worth) from 0.79 in 1989 to 0.85 in 2007. During this interval, the ratio of house value to net worth declined to below 0.65 in 2001 as a consequence of the surge in household financial wealth in the late 1990s. The homeownership rate, which had been fairly stable at 67 percent until 1995, started to increase in 1998, reached a high of 72.5 percent in 2004, and stayed high through 2007 (at 72.2 percent). Gabriel and Rosenthal (2005) show that demographic changes could explain a large part of this increase in the homeownership rate, although the relaxation of mortgage underwriting standards presumably also played a role.

Over the last 20 years, the most striking changes in the household balance sheet are related to stock ownership. Both the rate of participation in the stock market, and the average portfolio share of stocks increased dramatically over this period. The percentage of the households owning stocks (directly, or indirectly via retirement accounts and defined contribution pensions) was only 32.0 in 1989, increased to 51.4 in 2001 and further increased to 54.0 in 2007. Averaging across households, the level of stock holdings more than quadrupled over this period, increasing from a mean of \$31,900 in 1989 to \$130,600 in 2007. The median holding of stocks increased from zero in 1989, 1992, and 1995 to a modest \$2,500 in 2007, reflecting the fact that even in 2007 only slightly more than 50% of households owned positive amounts of stocks.

Table V Summary Statistics from the Survey of Consumer Finances, 1989-2007
(dollar amounts in thousands of constant 2007 dollars)

	1989		1992		1995		1998		2001		2004		2007	
	Mean	Med.	Mean	Med.	Mean	Med.	Mean	Med.	Mean	Med.	Mean	Med.	Mean	Med.
Age of Head	49.8	47	50.1	47	50.0	47	50.4	48	50.7	48	51.2	50.0	51.6	50.0
Transaction accounts	26.4	3.2	23.4	3.0	22.2	2.7	25.8	3.9	36.4	4.7	31.6	4.2	28.0	4.0
Stocks														
directly held	21.0	0	22.3	0	31.4	0	59.5	0	73.1	0	64.1	0	68.8	0
pensions & ret. a/c	6.6	0	12.2	0	17.8	0	33.1	0	46.4	0	42.9	0	49.6	0
other	4.2	0	2.3	0	4.5	0	8.9	0	17.2	0	11.2	0	12.2	0
Total	31.9	0	36.8	0	53.7	0	101.9	0.3	136.8	0.9	118.3	2.7	130.6	2.5
Financial assets	139.2	19.2	147.0	18.9	176.7	23.7	256.5	39.7	325.8	44.2	230.7	29.9	246.2	33.6
Primary residence ^{a/}	180.0	120.7	165.6	123.0	163.2	124.4	184.8	128.6	222.3	152.4	282.9	183.4	315.5	200.4
Mortgage outstanding ^{a/}	48.2	19.3	54.3	24.9	57.7	29.7	66.5	38.2	73.9	44.4	98.4	63.7	109.4	69.8
Net worth	347.3	100.8	332.9	95.0	359.1	105.6	470.7	133.5	595.3	151.9	551.3	136.8	620.4	154.0
stocks														
financial assets	0.093	0	0.121	0	0.144	0	0.195	0.054	0.212	0.070	0.271	0.170	0.265	0.146
house value ^{a/}														
net worth	1.164	0.788	1.614	0.792	2.816	0.754	1.383	0.650	1.142	0.646	1.656	0.844	1.792	0.852
Homeownership rate	67.1		67.1		66.6		69.2		70.5		72.5		72.2	
Ownership of stocks	32.0		37.0		39.9		49.4		51.4		53.9		54.0	
No. of obs.	2,806		3,442		3,958		3,911		4,073		4,129		4,068	

Note: The sample is limited to the households with the head 24 years or older, and with non-negative net worth. Nominal values are adjusted for inflation using the CPI-U deflator of each survey year provided in the Federal Reserve's web page of the SCF. Summary statistics are calculated with the analysis weight provided with the SCF.

^{a/} Mean and median are calculated only for homeowners.

Figure 1 plots the rate of homeownership as a function of the age of the household head. Panel (a) reports the cross-section view; i.e., each line represents a different wave of the survey and reports the relationship between homeownership rate and age in the cross-section of data for that wave. Panel (b) presents the patterns of homeownership by cohort for selected cohorts.


Each line follows a given birth cohort over the subsequent waves of the survey.¹⁶ Despite the considerable increase in aggregate rate of homeownership between 1998 and 2007, the basic life-cycle pattern of the homeownership rate has been remarkably stable. That is, the homeownership rate increases rapidly as a function of age for young households, continues to increase until the household reaches the mid-fifties, remains at a high level as the household reaches the mid-seventies, and declines only modestly for households over the age of 79. Cohort plots appear similar to cross-sectional plots, indicating there is not much variation in the age profiles of homeownership across different cohorts.

Figure 2 reports the median value, among homeowners, of h (ratio of house value to net worth) against age of household head. Considering the cross-section of households in any of the seven waves of the survey, the median value of h is greater than unity for the youngest households, and drops to 0.5 or 0.6 for households in their mid-fifties. For the second half of the life-cycle, the median value of h is stable. A comparison of later waves of the survey against the earlier waves (for example, 2004 and 2007 versus 1989 and 1992) indicates that young households in 2004 and 2007 have substantially greater ratios of house value to net worth than the previous generation. For older cohorts, the increase in h between 2001 and 2004 is less pronounced and the ratio exhibits similar patterns throughout time


¹⁶ We follow standard practice and group households into 21 three-year cohorts by the age of the household head. Each cohort is indicated by its 3-year birth year interval. Most cohorts are present in all seven waves of SCF. For younger and older cohorts, the year they enter or exit the survey is indicated in the figures.

Figure 1 Homeownership Rate, 1989-2007

(a) Cross-Section View



(b) Cohort View



and across different cohorts. The decline in h over the working life of the household is easy to understand, given that household net worth rises substantially over the working years, while the demand for housing services is relatively constant. Young homeowners are highly leveraged and typically own homes valued at two to three times their net worth. As the household accumulates financial assets over the course of the life cycle, h declines steadily to about 0.6 in their late 50s, and remains at that level past retirement age. It is only at a very advanced age that homeowners seem to run down their financial assets faster than they downsize their home, resulting in a very modest increase in h after age 76.

Figure 3 plots the stock-ownership rate as a function of age of household head. In both the cross-section and cohort views, the stockownership rate peaks in the mid-50s and declines thereafter. Over the 18 year time span covered by the survey, the stock-ownership rate at any age has risen substantially. The cohort view indicates that the increase in stockownership rates over time is especially dramatic for cohorts born after 1933-35. Each subsequent cohort invests in the stock market at a much higher rate than the previous cohorts at the same age.

Taken together, Figures 2 and 3 may point to a relationship between housing investment and stock investment. For younger cohorts (1969-71 and 1972-74 birth cohorts), the value of h increased most between 2001 and 2004 (Figure 2 panel (b)). At the same time, these cohorts decreased stock-market participation rate (Figure 3 panel (b)). Older cohorts did not decrease their stock-market participation rate between these two years, as the changes in the value of h for these cohorts were small.


Figures 4 and 5 plot the mean and median portfolio shares of stocks, s , respectively, by year and by cohort. Mirroring the increase in the rate of stockownership, these eight plots indicate that the share of financial assets invested in stocks rose throughout the sample period with the largest increases concentrated among the younger cohorts. For cohorts born after 1933-35, each subsequent cohort

almost always invests a higher fraction of assets in stocks than the previous cohort at the same age.

For the older cohorts, s has also grown, although at a much slower pace.

Figure 2 Median House-to-Net-Worth Ratio, 1989-2007

(a) Cross-Section View



(b) Cohort View



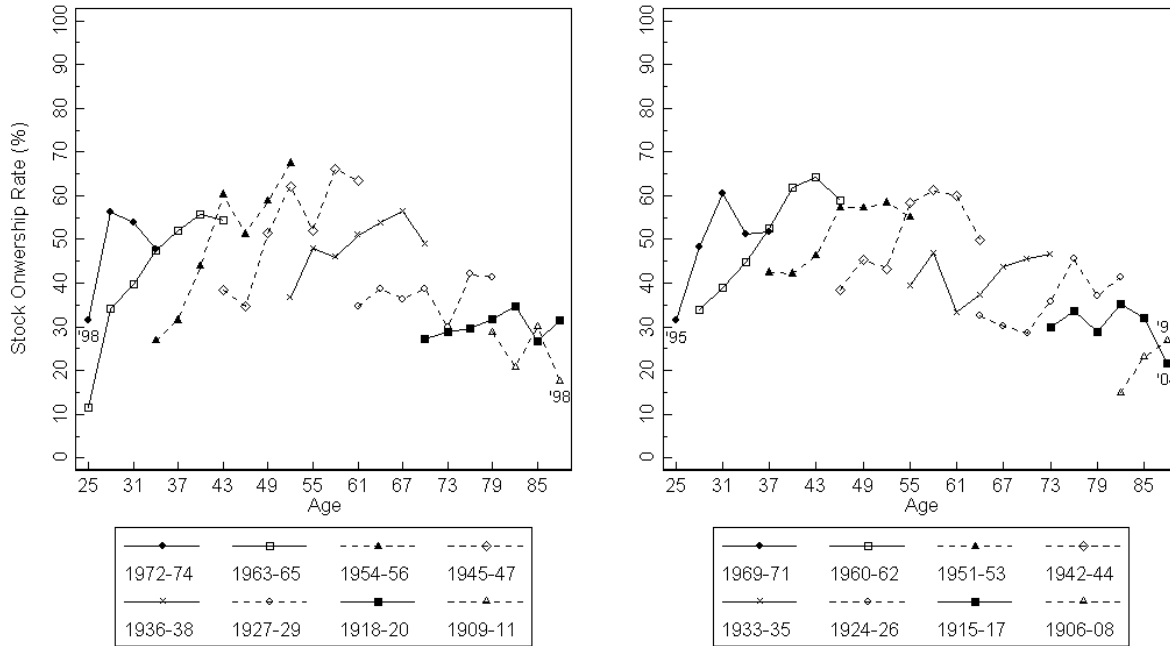


Figure 3 Stock-Ownership Rate, 1989-2007

(a) Cross-Section View



(b) Cohort View



The housing model implies that the variation in h over the lifecycle induces the hump-shaped age profile of the share of stock investment in financial assets. This hump-shaped age profile is most pronounced in 2004 and 2007 in Figures 4 and 5, when house values increased relative to other financial assets for many households. In addition, the cohort profiles of the younger cohorts rose parallel to each other (Figures 4(b) and 5(b)), indicating that each cohort has raised the portfolio share of stocks at the same pace,¹⁷ demonstrating that the stock market boom in the late 1990s creates the steep increases in cohort profiles. Of course other interpretations are possible. For example, a combination of age and cohort effects would create this pattern, i.e., each cohort holds more equity at any given age than an earlier cohort, and all individuals increase their holdings of stocks as they age, regardless of their birth cohort. Similarly, a combination of age and time effects could explain the

¹⁷ Ameriks and Zeldes (2000) offer interpretations for three patterns of cross-section and cohort variations in portfolio shares. For example, the pattern of s (in Figure 4) corresponds to Exhibit 1 for the younger cohorts and Exhibit 3 for the older cohorts in their paper.


pattern. A time effect (e.g., the stock market boom) causes all individuals to raise their portfolio share of stocks every period, while an age effect leads them to increase stock investment as they get older. Note, however, that both of these alternative explanations require an age effect on stock holding of the form that the share of the portfolio in stocks *increases* with age. Lifecycle portfolio models that focus on the role of human capital typically predict that households should decrease the share of the portfolio in riskier assets with age, since older households have less opportunity to offset investment losses with increased labor supply.

Figure 6 plots the mean and median of s against h .¹⁸ Disregarding the “dent” around $h = 1$, for the moment, Figure 6 indicates that homeowners with higher values of h invest a smaller proportion

Figure 4 Mean Portfolio Share of Stock Investment, 1989-2007

¹⁸ Households with the value of h between 0.05 and 1.05 are grouped into 0.1 intervals of h . For example, those with $0.05 \leq h < 0.15$ was grouped as $h = 0.1$, $0.15 \leq h < 0.25$ as $h = 0.2$, and so on. For larger values of h , the categories are coarser because the numbers of observations are small; we classify $1.05 \leq h < 1.2$ as $h = 1.125$, $1.2 \leq h < 1.45$ as $h = 1.325$, $1.45 \leq h < 2.05$ as $h = 1.75$, $2.05 \leq h < 2.95$ as $h = 2.5$, and $2.95 \leq h$ as $h = 3$. For very small values of h ($0 < h < 0.05$), the value of h is set to 0.001 to draw figures. This classification keeps the number of observations in each category roughly equal.

(a) Cross-Section View



(b) Cohort View





Figure 5 Median Portfolio Share of Stock Investment, 1989-2007

(a) Cross-Section View



(b) Cohort View





Figure 6 Portfolio Share of Stock Investment versus House-to-Net-Worth Ratio, 1989-2007

(a) Mean



(b) Median


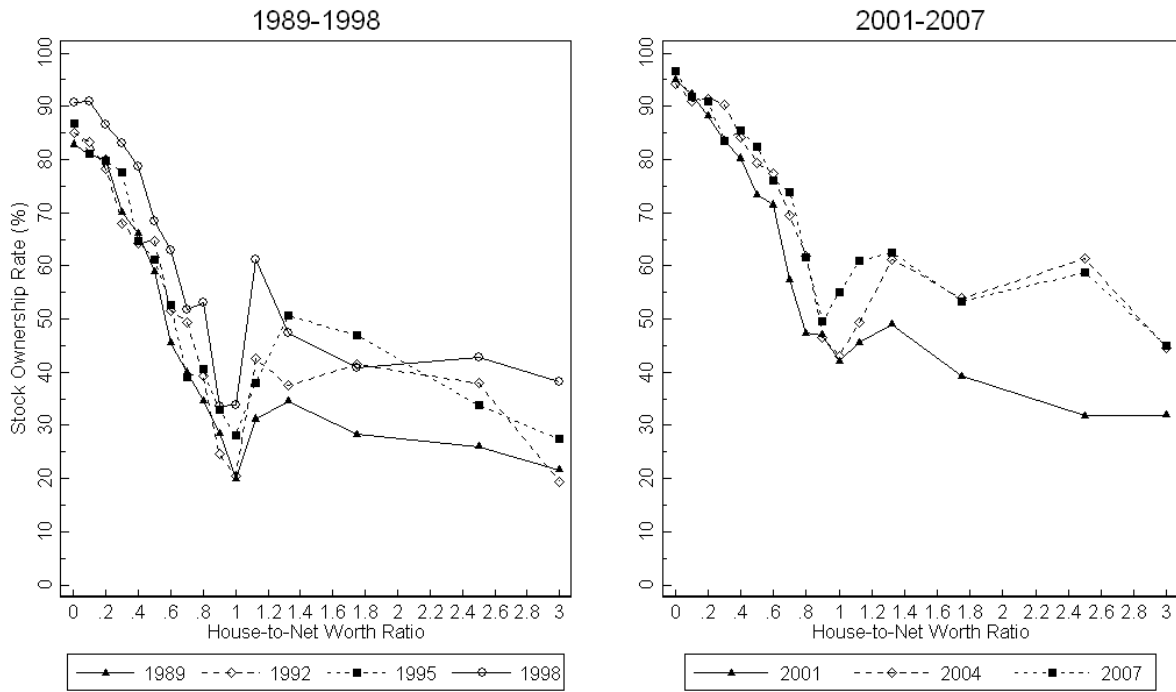


Figure 7 Stockownership Rate versus House-to-Net-Worth Ratio, 1989-2007



of financial assets in stocks. This inverse relationship between the portfolio share in stocks and the ratio of house value to net worth is also strongly articulated in Figure 6(b), which plots the median value of s against h . A plot of the stock-ownership rate (i.e., the percentage of households that own nonzero amounts of stocks) against h (Figure 7) closely mirrors the plot of the average portfolio share in stocks against h (Figure 6). Homeowners with high values of h are less likely to participate in the stock market, and the participation rate increases as h declines.¹⁹ In comparison, if we drop those not participating in the stock market from the sample, and plot the mean and median of equity share, s , against h for stock market participants (Figure 8), we find a very mild inverse relationship between s and h . Thus most of the increase in the average (over all households) portfolio share in stocks as h declines over the lifecycle is coming from an increase in the stock market participation rate, rather than an increase in the equity share of the average participant.

Notice also that the “dent” in the relationship between s and h around $h = 1$ is very pronounced in the plot of the stock market participation rate against h (Figure 7) but nearly absent in the plot of average equity share conditional on stock market participation (Figure 8). The overall sample includes a substantial number of households that report particularly simple balance sheets; these households (typically older, lower net worth households) own their homes with small or no mortgage but hold very little in the form of financial assets.²⁰ The portfolio behavior of these households may not be consistent with the optimal portfolios generated from the housing model. That is, given the assumptions on expected returns and the covariance matrix of returns, the housing model


¹⁹ Attanasio, Banks, and Tanner (2002) and Vissing-Jorgenson (2002) analyze the importance of non-participation in the stock market for asset pricing. Vissing-Jorgenson estimates that a small fixed cost of participation and transaction is sufficient to explain non-participation of the majority of non-stockholders.

²⁰ Households with the value of h near 1 have a substantially smaller amount of financial assets compared to those with lower or higher values of h . For example, in 2004 the median value of financial assets for households with $h = 0.9$ and 1 is \$16,288 and \$8,000, respectively, compared to \$40,204 and \$22,000 for $h = 0.8$ and 1.125, respectively. Similarly, the median age for those with $h = 0.9$ and 1 is 58 and 52 in 2004, respectively, compared to 55 and 48 for $h = 0.8$ and 1.125, respectively.


Figure 8 Portfolio Share of Stocks vs. House-to-Net-Worth Ratio (Stockholders Only),

1989-2007

(a) Mean



(b) Median



implies that a household with net worth equal to the value of its home ($h = 1$), should borrow against the house and invest the proceeds in the stock market. At low levels of risk aversion, the optimizing household should borrow the maximum (100% of the house value) and invest all of the proceeds in the stock market. Even moderately risk averse households would optimally borrow against the house to some extent (10% to 30% of the value of the house, depending on the degree of risk aversion) in order to hold some equities. Thus the existence of a sizeable subsample of households that hold essentially all of their net worth in the form of a mortgage-free home is not consistent with the model. Note that the result that an optimal portfolio will always include a nonnegative amount of stocks arises in any model which assumes frictionless transaction in the stock market, not just the housing model. If participation in the stock market is costless, even the most risk averse household would hold at least \$1 in equities.


Nonparticipation in the stock market also plays an important role in the relationship between equity share and age. In Figures 9 and 10, we plot the mean and median of s against age after eliminating nonparticipants from the sample. In contrast to the age profiles of s for the entire sample (Figures 4 and 5), which exhibit a hump-shape pattern, the portfolio share of stocks (conditional on stockownership) is mildly increasing with age.

To demonstrate that the stock investment over life cycle is driven by the house-value-to-net worth ratio, we consider two more figures. In Figure 11, we plot the age profile of average portfolio share in equities for *non-homeowners*. A comparison of Figure 11 (for non-homeowners) to Figure 4 (for homeowners) indicates that the age profile of the equity share is mildly decreasing or flat with age for non-homeowners in contrast to the hump shaped age profile for homeowners.²¹ Although younger cohorts among non-homeowners increase the share of stock investment as they age, the rate of increase

²¹ The median of s for non-homeowners is zero for all ages throughout all years. Thus we only plot the mean.

Figure 9 Mean Portfolio Share of Stock Investment (Stockholders Only), 1989-2007

(a) Cross-Section View



(b) Cohort View








Figure 10 Median Portfolio Share of Stock Investment (Stockholders Only), 1989-2007

(a) Cross-Section View



(b) Cohort View

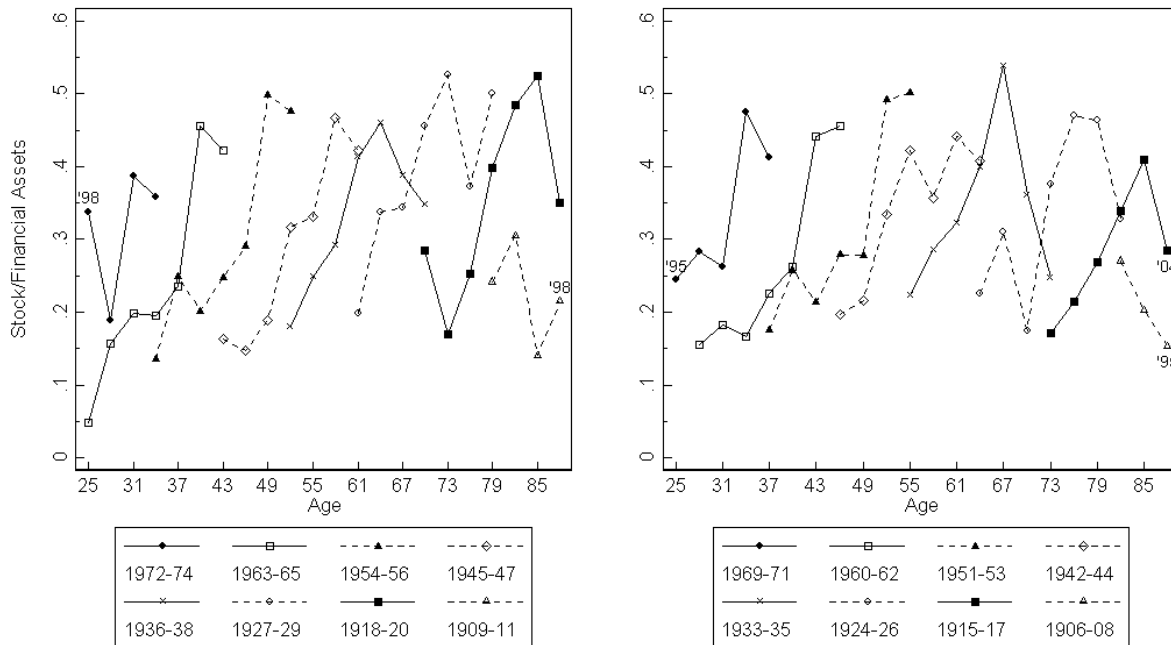



Figure 11 Mean Portfolio Share of Stock Investment (Non-Homeowners Only), 1989-2007

(a) Cross-Section View



(b) Cohort View

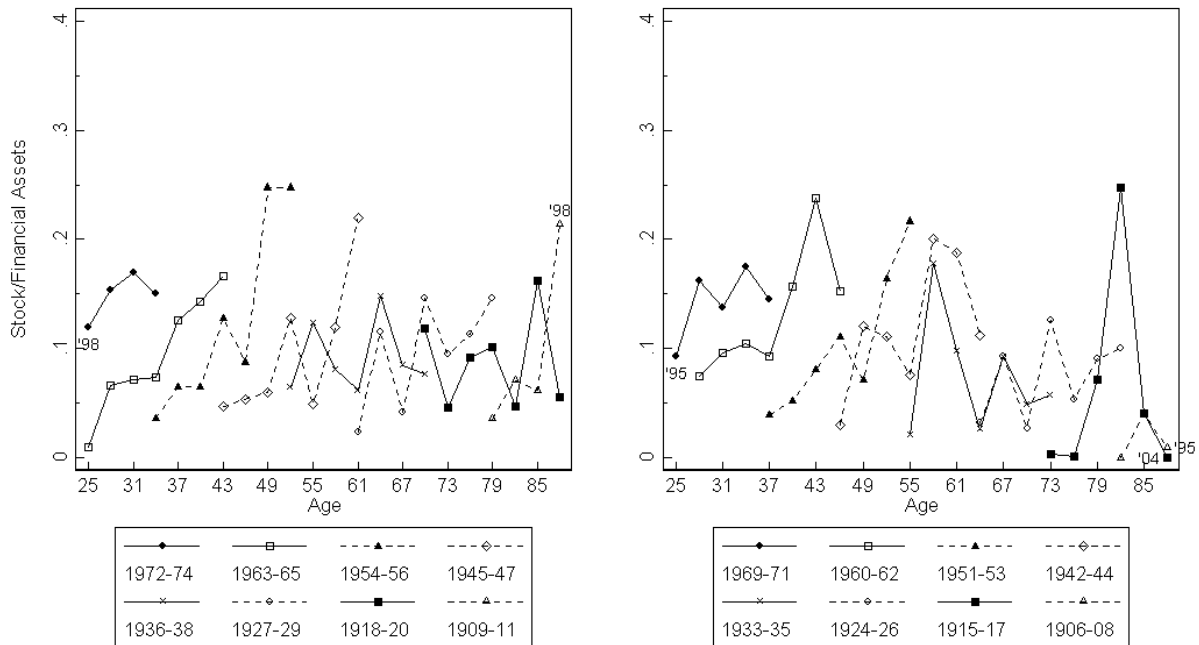




Figure 12 Portfolio Share of Stocks vs. LTV Ratio, 1989-2007

(a) Mean



(b) Median



is slower than that of homeowners. Figure 12 plots the share of stocks in portfolios of financial assets against loan-to-value (LTV) ratio among homeowners. In earlier years, there is very little relationship between the two. In more recent years, particularly in 1998 and 2001, the stock investment is negatively related to the LTV ratio. However, the association between the value of s and LTV ratio is weak compared to the strong negative relationship between the values of s and h shown in Figure 6. The housing model thus seems to offer a plausible and consistent explanation of the lifecycle pattern of stock holding.

Section IV: Statistical tests of the role of the housing state variable

According to the housing model, the hump-shaped age profile of stock holding is a result of the decline in the ratio of house value to net worth over the lifecycle. In regressions which fail to control for the housing state variable, however, the homeowner's age will act as a proxy for the omitted state variable even if "age" itself is not a determinant of the portfolio. If the lifecycle pattern of stock holding is induced by variation in the housing state variable, rather than by age itself, we would expect the predictive value of variables reflecting the homeowner's age to disappear when the housing state variable – the ratio of house-value-to-net-worth -- is added as a control in a regression of the ratio of stocks-to-financial-assets on age. Using data from the SCF, we estimate regressions of the stock-to-financial-assets ratio on twenty categories of age dummies (age 24-26 is the omitted category). An important advantage of the SCF is that the survey asks questions about the respondent's attitudes toward risk, financial planning, and saving behavior. To elicit a measure of risk tolerance, the survey asks whether or not the respondent is "willing to take bigger financial risks for higher returns." In addition, the survey asks whether or not the respondent "is willing to take any financial risks." Since the housing model implies that stock holdings will be a decreasing function of the housing state

variable for a given degree of risk aversion, the regressions rely on these two indicator variables (willingness to take above-average risk for above-average returns, and not willing to take any financial risks) to control for self-reported attitudes towards risk. We also include other controls that may capture household characteristics related to stock investment, such as education, race, marital status, gender, and log of labor income as well as the respondent's planning horizon, attitudes toward saving, and six dummy variables for the year of survey.²²

Table VI Regression Results with and without House-to-Net-Worth Ratio

	OLS				Quantile Regression (q = 0.5)			
	coeff.	s.e.	coeff.	s.e.	coeff.	s.e.	coeff.	s.e.
	(1)		(2)		(3)		(4)	
h			-0.024**	(0.003)			-0.019**	(0.001)
age 27-29	0.007	(0.022)	0.004	(0.022)	0.001	(0.012)	-0.002	(0.012)
age 30-32	0.015	(0.020)	0.011	(0.020)	0.009	(0.012)	-0.004	(0.011)
age 33-35	-0.001	(0.019)	-0.008	(0.019)	-0.010	(0.012)	-0.016	(0.011)
age 36-38	0.027	(0.019)	0.016	(0.019)	0.012	(0.011)	-0.002	(0.011)
age 39-41	0.053**	(0.020)	0.037	(0.020)	0.025*	(0.011)	0.005	(0.010)
age 42-44	0.038*	(0.019)	0.022	(0.019)	0.023*	(0.010)	0.001	(0.010)
age 45-47	0.044*	(0.019)	0.025	(0.019)	0.030**	(0.011)	0.007	(0.010)
age 48-50	0.058**	(0.019)	0.036	(0.019)	0.028**	(0.011)	0.003	(0.010)
age 51-53	0.076**	(0.020)	0.053**	(0.020)	0.049**	(0.011)	0.022*	(0.011)
age 54-56	0.077**	(0.020)	0.052**	(0.020)	0.050**	(0.013)	0.023*	(0.011)
age 57-59	0.108**	(0.020)	0.081**	(0.020)	0.065**	(0.012)	0.037**	(0.011)
age 60-62	0.071**	(0.020)	0.042*	(0.020)	0.041**	(0.011)	0.012	(0.011)
age 63-65	0.092**	(0.021)	0.063**	(0.021)	0.038**	(0.012)	0.009	(0.011)
age 66-68	0.094**	(0.020)	0.064**	(0.021)	0.045**	(0.012)	0.015	(0.011)
age 69-71	0.060**	(0.021)	0.029	(0.021)	0.030**	(0.012)	-0.000	(0.011)
age 72-74	0.063**	(0.021)	0.032	(0.021)	0.031**	(0.012)	-0.000	(0.011)
age 75-77	0.084**	(0.022)	0.053*	(0.023)	0.029*	(0.012)	-0.001	(0.012)
age 78-80	0.081**	(0.023)	0.048*	(0.023)	0.037**	(0.013)	0.006	(0.012)
age 81-83	0.064**	(0.026)	0.030	(0.026)	0.032*	(0.013)	-0.002	(0.014)
age 84-86	0.045*	(0.026)	0.012	(0.026)	0.020	(0.014)	-0.015	(0.014)
age 87-89	0.020	(0.032)	-0.014	(0.032)	0.018	(0.016)	-0.015	(0.016)

* significant at the 5% level

** significant at the 1% level

²² We limit our sample to homeowners only. As some households have extreme values of the house-to-net worth ratio (over 1,000), we trim the top 1% of the value of h and retain the sample with the value of h smaller than 6.82.

The number of observations is 20,007. All five imputations are used for estimation, and standard errors take into account sample variation across different imputations using Stata's `micombine` command. Standard errors of OLS estimates are Huber-White robust standard errors. Regressions also control for three categories of educational attainment (high-school dropout, some college, and college), two categories of race (black and "other race"), three categories of marital status (single, divorced/separated, and widowed), four categories of spousal education, female headship, two categories of degree of risk aversion, two categories of planning horizon, an indicator for not saving regularly, log of labor income, and six dummy variables for the year of survey.

Column (1) of Table VI reports the results of the ordinary least squares (OLS) estimation of the age effects when the housing state variable is omitted. The coefficients on the age dummies increase with age, peaking at age 57-59, and then decline as the household head becomes older, confirming the hump-shaped age profile. The estimates are statistically significant at the 5% level for all 15 categories between age 39 and 83. In column (2), we include the house-to-net-worth ratio as an additional control. As the housing model predicts, the coefficient on h is negative and statistically highly significant. Controlling for the house-to-net-worth ratio, furthermore, the coefficients on age dummies become substantially smaller, by 25% to 50%, making the age profile less pronounced. In addition, most estimates become statistically insignificantly different from zero; only eight categories (age 51-68, and 75-80) remain significant at the 5% level.


Columns (3) and (4) repeat the same specifications with quantile regressions estimated at the median. As in the OLS results, a hump-shaped age profile is detected with the peak at age 57-59, although the estimates are smaller than OLS estimates in magnitude. In the quantile regression, the impact of the inclusion of h on the age coefficients is larger, even though the estimate of the direct effect of h is smaller in magnitude, compared to the OLS results. The coefficients on age dummies

are reduced by 50 to 90 percent compared to the estimates without h and all but three age categories lose statistical significance.

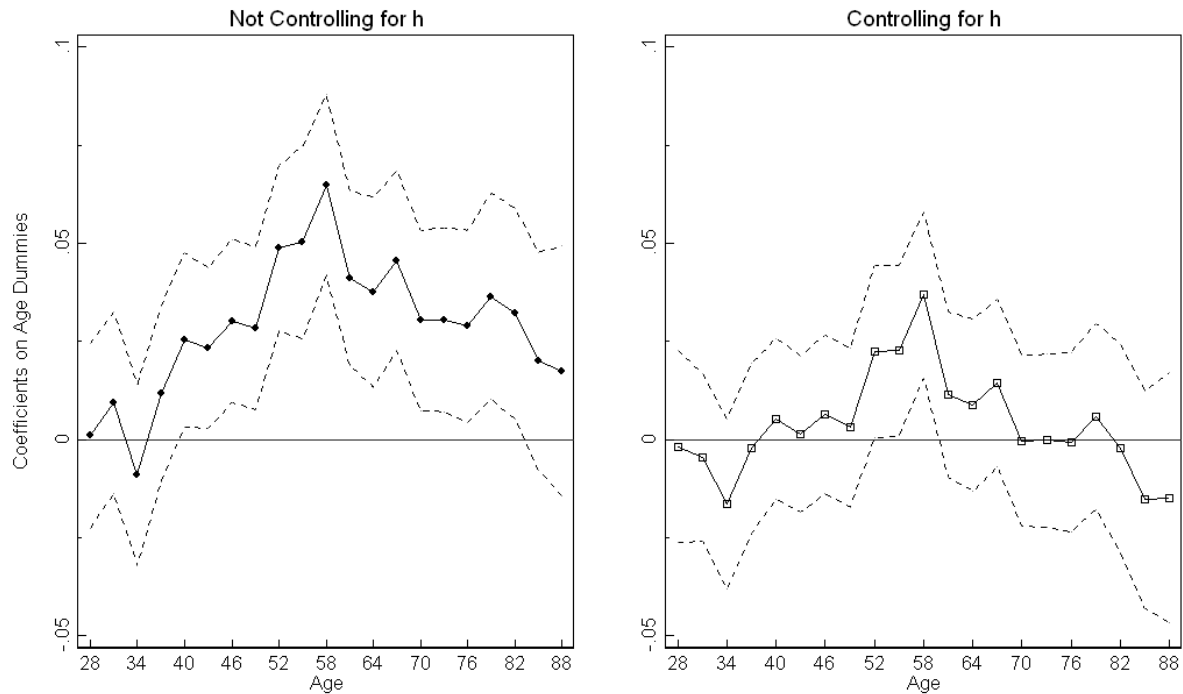
Figure 13 plots the estimated age coefficients, with and without controlling for the housing state variable, h . Solid lines represent the coefficient estimates (panel (a) for OLS and panel (b) for quantile regression) and dotted lines indicate the 95% confidence interval. Inclusion of h shifts down the age profile substantially, and also makes the age profile less steep.

Given that the left hand side variable, and therefore the errors, are right skewed, the OLS estimates will be more sensitive than the quantile regression estimates to the relatively small number of households who hold a large fraction of their portfolio in stocks, which explains why the quantile regression estimates are both somewhat smaller in magnitude and more precisely estimated than the OLS estimates. Nevertheless, both OLS and quantile regression deliver similar estimates of the coefficient on the state variable, h . The OLS estimate is -0.024 with a standard error of 0.003 ; the quantile regression estimate is -0.019 with a standard error of 0.001 . In order to interpret the magnitude of the coefficient on the housing state variable, consider a household that experiences a decline in the ratio of house value to net worth from 3.5 as a young household to $.5$ as a retired household. A decline in the state variable from 3.5 to $.5$ would imply an increase in fraction of financial assets held in the form of stocks of 5.7% (based on the quantile regression point estimate) to 7.2% (based on the OLS estimate). While a 5.7 to 7.2% increase in the fraction of financial assets held in the form of stocks may seem, at first glance, to indicate that the housing state variable has a fairly modest quantitative effect on the composition of the portfolio, it is important to interpret the magnitude of the effect relative to the average level of stock holdings for a typical household in the survey. In the cross section of households represented in the 1989 sample, the variable representing stock holdings as a fraction of total financial assets had a mean of 9.3% and a median of 0% . By the time the 2007 survey

Figure 13 Coefficient Estimates on Dummy Variables on Age
 (a) Estimates from OLS



(b) Estimates from Quantile Regressions at the Median



was conducted, stock holding had increased substantially. However, even in 2007, the mean value of stock holdings as a fraction of total financial assets was only 26.5% (median value was 14.6%). Considered in the context of the average or typical ratio of stock holdings to financial assets in the cross section, the estimated increase in stocks as a fraction of financial assets of 5.7 to 7.2% is a quantitatively important effect.

Section V: Conclusions

The paper models the dual role of housing as both a consumption good and an asset in the household portfolio in a continuous time framework. The model allows for variation in the relative price of housing, both over time and across different regional housing markets. Nondurable consumption and the holdings of financial assets are assumed to be frictionlessly adjustable; in contrast, adjustment of the quantity of housing is assumed to be subject to a nonconvex adjustment cost. While the adjustment cost on housing greatly complicates some aspects of the optimization problem, it actually

simplifies analysis of the portfolio allocation problem by making the household's decision process recursive. That is, at every moment the household considers whether or not it is worthwhile to pay the adjustment cost and reoptimize over the current holding of housing. Having decided that it is not optimal to change the quantity of housing immediately, the household then takes the current quantity of housing as a state variable, and determines the optimal holding of financial assets conditional of the value of the state variable. If housing price risk is uncorrelated with the returns to financial assets (an assumption that is consistent with the data), the household's optimal portfolio of financial assets will be mean-variance efficient.

The model focuses on the role of housing as providing collateral by imposing the constraint that the household may borrow only in the form of a mortgage, and, further, that the value of the mortgage cannot exceed the value of the house. Holdings of financial assets are constrained to be nonnegative. For a range of assumptions concerning the stochastic structure of asset returns, we calculate the household's optimal portfolio of financial assets, conditional of the value of the housing state variable and the household's level of risk aversion, and find that the collateral and nonnegativity constraints are often binding.

Since the state variable – the ratio of house value to net worth – varies over the lifecycle, the constrained mean-variance frontier available to the household, and therefore the optimal portfolio, also varies over the lifecycle. Based on numerical optimization of the constrained optimal portfolios, we find that, holding constant the household's relative risk aversion, the fraction of financial assets held in the form of stocks generally increases over the lifecycle. Empirical work using data from the SCF indicates that the effect of the housing state variable on the fraction of financial assets held in the form of stocks is quantitatively and statistically significant, and generally consistent with the predictions of the model.

References

- Ameriks, John and Stephen P. Zeldes, 2000, "How Do Household Portfolio Shares Vary with Age?" Unpublished manuscript, Columbia University.
- Attanasio, Orazio P., James Banks, and Sarah Tanner, 2002, "Asset Holding and Consumption Volatility," *Journal of Political Economy* 110, 771-792.
- Buiter, Willem H., 2008, "Housing Wealth Isn't Wealth." *NBER Working Paper 14204*, Cambridge, MA.
- Cauley, Stephen Day, Andrey D. Pavlov, and Eduardo S. Schwartz, 2007, "Homeownership as a Constraint on Asset Allocation," *Journal of Real Estate Finance and Economics*, 34, 283-311.
- Cocco, Joao, 2005, "Portfolio Choice in the Presence of Housing," *Journal of Financial Studies*, 18, 535-567.
- Faig, Miguel, and Pauline Shum, 2002, "Portfolio Choice in the Presence of Personal Illiquid Projects," *Journal of Finance*, 57, 303-328.
- Federal Reserve Board, 2008, SAS Program Used to Create Summary Variables: Detailed Definitions, Washington, DC.
- Flavin, Marjorie, 2009, "Housing, Adjustment Costs, and Endogenous Risk Aversion", working paper, UCSD.
- Flavin, Marjorie, and Shinobu Nakagawa, 2008, "A Model of Housing in the Presence of Adjustment Costs: A Structural Interpretation of Habit Persistence," *American Economic Review*, 98, 474-495.
- Flavin, Marjorie, and Takashi Yamashita, 2002, "Owner-Occupied Housing and the Composition of the Household Portfolio," *American Economics Review*, 92, 345-362.
- Gabriel, Stuart A., and Stuart S. Rosenthal, 2005, "Homeownership in the 1980s and 1990s: Aggregate Trends and Racial Gaps," *Journal of Urban Economics*, 57, 101-127.
- Gomes, Francisco, and Alexander Michaelides, 2005, "Optimal Life-Cycle Asset Allocation: Understanding the Empirical Evidence," *Journal of Finance*, 60, 869-904.
- Grossman, Sanford J., and Guy Laroque, 1990, "Asset Pricing and Optimal Portfolio Choice in the Presence of Illiquid Durable Consumption Goods," *Econometrica*, 58, 25-51.
- Hu, Xiaoqing, 2005, "Portfolio Choices for Homeowners," *Journal of Urban Economics*, 58, 114-136.
- Hurst, Erik, and Annamaria Lusardi, 2004, "Liquidity Constraints, Household Wealth, and Entrepreneurship," *Journal of Political Economy*, 112, 319-347.
- Vissing-Jorgensen, Annette, 2002, "Limited Asset Market Participation and the Elasticity of Intertemporal Substitution," *Journal of Political Economy*, 110, 825-853.
- Yao, Rui, and Harold H. Zhang, 2005, "Optimal Consumption and Portfolio choices with Risky Housing and Borrowing Constraints," *Journal of Financial Studies*, 18, 197-239.