

Portfolio Choice in the Presence of Housing

João F. Cocco

London Business School

I show that investment in housing plays a crucial role in explaining the patterns of cross-sectional variation in the composition of wealth and the level of stockholdings observed in portfolio composition data. Due to investment in housing, younger and poorer investors have limited financial wealth to invest in stocks, which reduces the benefits of equity market participation. House price risk crowds out stockholdings, and this crowding out effect is larger for low financial net-worth. In the model as in the data leverage is positively correlated with stockholdings.

Owner-occupied housing is the single most important asset in many investors' portfolios. This article addresses the question of how the investment in housing affects the composition of an investor's portfolio. In particular, do house price risk and the illiquid nature of the housing investment lead investors to reduce their exposure to stocks?

For younger and poorer households, the investment in housing is often financed through a mortgage contract to create a leveraged position in residential real estate. Heaton and Lucas (2000a) find that for stocks relative to financial assets, "a higher mortgage leads to higher stock holdings, suggesting that some stocks are indirectly financed via mortgage debt."¹ Why might a (risky) leveraged position in residential real estate be associated with a larger investment in risky financial assets?

To answer these questions, I solve a model of the optimal portfolio and consumption decisions of an investor who is endowed with nontradable human capital. Human capital generates dividends in the form of labor income.² Housing ownership also generates dividends, but in the form of

I would like to thank the editor, John Heaton, and two anonymous referees for comments that greatly improved this paper. This paper is a substantially revised version of Chapter 4 of my Harvard Ph.D. Thesis. I am grateful to my thesis committee—John Y. Campbell, David I. Laibson, and N. Gregory Mankiw—for guidance. I have also benefited from comments by Rui Albuquerque, Alberto Alesina, Francisco Gomes, Paula Lopes, Pascal Maenhout, Daniele Paserman, Andrei Shleifer, Raman Uppal, Luis Viceira, David Weil, and seminar participants at Cemfi, Columbia, Harvard, Insead, Kellogg, London Business School, the NBER Finance Lunch, Pompeu Fabra, Princeton, Virginia, and Wharton. Financial support from the *Banco de Portugal* and *Fundação para a Ciência e Tecnologia* is gratefully acknowledged. Address correspondence to: João F. Cocco, London Business School, Regent's Park, London NW1 4SA, UK, Tel.: (020) 7262-5050 or e-mail: jcocco@london.edu.

¹ In spite of the word "leads" in this sentence Heaton and Lucas make clear that their results are correlations and do not imply causality.

² If markets are complete so that future labor income can be capitalized and its risk insured, then human capital can simply be added to current wealth and plays no particular role. But moral hazard issues prevent investors from borrowing against future labor income, and insurance markets for labor income risk are not well-developed.

consumption services from which the investor derives utility. Thus, in my model, housing has a dual dimension, as an asset in a portfolio and also as a consumption good. The value of the house limits the degree of leverage of the investor's portfolio, and the investor allocates his financial savings between a risky (stocks) and a riskless financial asset.

Following the literature on limited stock market participation [Basak and Cuoco (1998), Luttmer (1999), Polkovnichenko (2000), Vissing-Jorgensen (2002)], I assume that there is a fixed cost of equity market participation, and ask how the housing investment affects investors' willingness to pay the fixed cost.³

The investment in housing may affect the composition of the investors' portfolio also because the price of residential real estate may be correlated with labor income shocks and stock returns. I estimate these correlations using PSID data, to find that aggregate income shocks are strongly positively correlated with house price shocks, but that these are uncorrelated with stock returns. These correlations are used to parameterize the model.

The results show that investment in housing plays a crucial role in explaining the patterns of cross-sectional variation in the composition of wealth and the level of stockholdings observed in portfolio composition data. Due to investment in housing, younger and poorer investors have limited financial wealth to invest in stocks, which reduces the benefits of equity market participation. Therefore, in the presence of housing, a lower fixed cost of equity market participation is needed to generate the empirically observed levels of participation.

I also find that house price risk crowds out stock holdings, both for high and low financial net-worth investors, but this crowding out effect is larger at low levels of financial net-worth. The crowding out effect of both house price risk and a minimum house size are important for explaining the level of stockholdings observed in the data.

Finally, the model provides an explanation as to why in cross-sectional data, leverage and investment in risky financial assets appear to be positively correlated. Due to the consumption dimension of housing, investors who have more human capital acquire more expensive houses and borrow more. At the same time, human capital, although risky, resembles Treasury bills more closely, inducing a tilt in the financial portfolio toward stocks, as in Heaton and Lucas (1997) and Viceira (2001).

There is a vast literature on portfolio selection,⁴ but most of this literature ignores housing. Exceptions include Grossman and Laroque

³ One well-documented feature of the data on portfolio composition is that many households, particularly poorer and younger ones, do not own stocks at all. This is inconsistent with simple frictionless models of portfolio choice, but may be explained if there is a fixed cost of equity market participation.

⁴ Merton (1971) and Samuelson (1969) are classical references.

(1991) who develop an asset allocation model with a single illiquid durable consumption good from which an infinitely lived investor derives utility, and Cuoco and Liu (2000) who consider instead a divisible durable good. Absent from their analyses are house price risk and nontradable income, which play a crucial role in this article. Flavin and Yamashita (2002) study the impact of the portfolio constraint imposed by the consumption demand for housing, which they call the “housing constraint,” on investors’ optimal holdings of financial assets. At each age, the ratio of housing to net-worth is taken to be equal to the value in PSID data, and mean-variance analysis is used to characterize optimal portfolios. Faig and Shum (2002) study portfolio choice in the presence of personal illiquid projects, which may be interpreted as housing.

The article closest to mine is that of Yao and Zhang (2004), who also study the effects of housing on the portfolio allocation of liquid wealth among stocks and bonds. They find that when investors are indifferent between renting and owning a house, they choose substantially different portfolio allocations when owning a house versus when renting housing services. When owning a house, investors substitute home equity for risky stocks, but hold a higher equity proportion in their liquid financial portfolio (bonds and stocks). Unlike Yao and Zhang (2004), I do not study the effects of the renting versus owning decision on portfolio allocation, but my model incorporates a fixed cost of equity market participation, and I study how the housing investment affects investors’ willingness to pay the fixed cost. This is important for explaining limited equity market participation.

The article is organized as follows. In Section 2, I present the set-up of the model. In Section 3, I use PSID data to parameterize the model. Section 4 presents the results. Section 5 compares these results to PSID data on portfolio composition. This comparison is important since it provides evidence of the strengths and limitations of the model. Section 6 discusses the model and results. The final section concludes the discussion.

1. The Model

1.1 Preferences

I model the asset and consumption choices of an investor who lives for T periods.⁵ In each period t , the investor needs to choose the size of house to own, H_t , and other nondurable goods consumption, C_t . The date t price

⁵ It is possible to extend the model to allow for uncertain life span in the manner of Hubbard, Skinner, and Zeldes (1994). Uncertain life span probably is one of the reasons why old households reduce their consumption of housing services only late in life, often precipitated by widowhood. Uncertain life span does not however explain why reverse annuity mortgages are not more widely used.

per unit of housing is denoted by P_t , such that a house of size $\bar{H}$ has price $P_t \bar{H}$ at date t . The size of the house should be interpreted broadly as reflecting not only the physical size but also its quality. The price of the consumption of other goods (the numeraire) is fixed and normalized to 1. The investor derives utility from both housing and nondurable goods, and after date T from bequeathing terminal wealth, W_{T+1} . Preferences are described by:

$$U_1 = E_1 \sum_{t=1}^T \beta^{t-1} \frac{(C_t^{1-\theta} H_t^\theta)^{1-\gamma}}{1-\gamma} + \beta^T \frac{W_{T+1}^{1-\gamma}}{1-\gamma}, \quad (1)$$

where β is the time discount factor, γ is the coefficient of relative risk aversion, and θ measures preference for housing relative to nondurable consumption goods.

1.2 Labor income risk

The investor works for the first K periods of his life ($K \leq T$) in which labor is supplied inelastically.⁶ At each date $t \leq K$ the investor receives a stochastic labor income stream $\tilde{Y}_t$, against which he cannot borrow. Let lower case letters denote the log of the variable, that is, $y_t \equiv \ln(Y_t)$. At investor i 's age t , labor income is exogenously given by:

$$\tilde{y}_{it} = \begin{cases} f(t, Z_{it}) + \tilde{u}_{it}, & \text{for } t \leq K, \\ f(t, Z_{it}), & \text{for } t > K, \end{cases} \quad (2)$$

where $f(t, Z_{it})$ is a deterministic function of age, t , and other individual characteristics, Z_{it} , and $\tilde{u}_{it}$ can be decomposed into an aggregate (η_t) and idiosyncratic components (ω_{it}) such that:

$$\tilde{u}_{it} = \tilde{\eta}_t + \tilde{\omega}_{it}. \quad (3)$$

I assume that idiosyncratic labor income risk is transitory so that $\tilde{\omega}_{it}$ is an i.i.d. normally distributed random variable with mean zero and variance σ_ω^2 . The aggregate shock, $\tilde{\eta}_t$, follows a first-order autoregressive [AR(1)] process:

$$\tilde{\eta}_t = \phi \eta_{t-1} + \tilde{\epsilon}_t, \quad (4)$$

where ϵ_t is an i.i.d. normally distributed random variable with mean zero and variance σ_ϵ^2 . Thus, prior to retirement, log income is the sum of a deterministic component that can be calibrated to capture the hump-shape of earnings over the life-cycle, and two random components, one transitory and one persistent. Log income in retirement is modeled as a

⁶ Bodie, Merton, and Samuelson (1991) have studied the effects of labor supply flexibility on portfolio choice.

deterministic function of age and other individual characteristics [$f(t, Z_{it})$] reflecting the fact that at this stage of life most of the uncertainty related to future labor income has been resolved.

1.3 Housing

As for owner-occupied housing, I assume a correspondence between the size of the house the investor owns and the consumption benefits that he derives from it.⁷ One important feature of housing that makes it different from liquid financial assets is its indivisibility. To model it, I assume that there is a minimum house size, $\bar{H}_{\min}$, so that:

$$H_t \geq \bar{H}_{\min}, \quad \forall t. \quad (5)$$

The model generates endogenous house trades. But in practice, there are other reasons for households to move, exogenous to the model and related to changes in employment or family specific shocks. To capture these, I assume that in each period t with probability π the household is forced to sell the house and buy another one (there is an involuntary move). With probability $(1 - \pi)$, the household is not forced to move, but may still do so if that is optimal. Introducing involuntary house trades in the model requires a state variable, $InvMove_t$ which takes the value of 1 if at t the household is in a state where it is forced to move, and zero otherwise.⁸

Frequently, the price of housing in a given region is affected by labor income shocks in the same region. I capture this by assuming that cyclical fluctuations in house prices are correlated with labor income shocks. Let p_t denote the date t log price of one unit of housing, and $p'_t \equiv p_t - bt$ the detrended log price of housing. I assume that cyclical fluctuations in house prices are perfectly positively correlated with aggregate labor income shocks,⁹ and imperfectly correlated with temporary labor income shocks so that:

$$\eta_t = \kappa_\eta \tilde{p}'_t, \quad (6)$$

$$\omega_{it} = \kappa_\omega \tilde{p}'_t + \zeta_{it}, \quad (7)$$

⁷ Thus, as Grossman and Laroque (1991), I ignore rental markets. Rental markets allow investors to separate the consumption and investment dimensions of housing. In this article, I assume that there are market frictions that make buying a strictly preferred alternative. Possible market frictions include taxes, transaction costs, and moral hazard. Hu (2002) and Yao and Zhang (2004) study the choice of renting versus owning and its impact on the financial portfolio.

⁸ I would like to thank an anonymous referee for suggesting the incorporation of a minimum house size and involuntary house trades in the model.

⁹ The assumption of correlation equal to 1 greatly simplifies the solution of the problem since it avoids the introduction of one additional state variable. In the parameterization section, I use PSID data to estimate the correlation between p'_t and η_t . This allows us to assess how reasonable the assumption of perfect correlation is.

where κ_η and κ_ω are regression coefficients, and ζ_{it} is a normally distributed variable with zero mean and variance σ_ζ^2 . By varying κ_ω , I can evaluate the effects of the correlation between labor income and house price shocks on the investor's portfolio.

To capture the illiquid nature of the housing investment, I assume that a house sale is associated with a monetary cost equal to a proportion λ of the house value. Annual maintenance costs are equal to a proportion δ of the house value.

1.4 Financial assets and credit markets

There are two financial assets. A riskless asset, called *Treasury bills*, with gross real return $\bar{R}_F$ and a risky asset, called *stocks*, with gross real return $\bar{R}_t$. The log return on the risky asset $\bar{R}_t$ is assumed to be:

$$\log(\bar{R}_t) = \mu + \tilde{r}_t, \quad (8)$$

where $\mu > 0$ is the expected log return and $\tilde{r}_t$, the innovation to log returns, is assumed to be distributed as $N(0, \sigma_t^2)$. I allow innovations in log returns to be correlated with innovations to aggregate income shocks (and house prices), and denote the corresponding coefficient of correlation by $\rho_{\epsilon, \tilde{r}}$. The dollar amount the investor has in bills and stocks at date t are denoted B_t and S_t , respectively. I assume that the investor cannot short-sell either of these assets such that:

$$S_t \geq 0, \quad B_t \geq 0, \quad \forall t. \quad (9)$$

These restrictions may be motivated by the costs associated with short positions. I allow for fixed costs of equity market participation: to have access to equity markets, the investor has to pay a one-time monetary fixed cost equal to F .

A third financial asset, which I call a *mortgage*, allows the investor to borrow against the value of the house, at a gross real fixed rate of $\bar{R}_D$. The dollar amount the investor owes in mortgages at date t is denoted D_t . I assume that the investor may borrow up to the value of the house minus a down payment, which is assumed to be a proportion (d) of the value of the house, such that:

$$D_t \leq (1 - d)P_t H_t, \quad \forall t. \quad (10)$$

I assume that the investor is allowed in every period to costlessly renegotiate the desired level of debt, as it is the case for a home line of credit.¹⁰

¹⁰ The possibility of (in every period) costlessly renegotiating the level of outstanding debt up to constraint (10) greatly simplifies the numerical solution of the problem since it avoids having the level of outstanding debt as a state variable. Note that in the presence of costless debt adjustments, households are more likely

1.5 The investor's optimization problem

The investor starts period $t > 1$ with liquid wealth (LW_t) given by:

$$LW_t = \tilde{R}_t S_{t-1} + \bar{R}_f B_{t-1} - \bar{R}_D D_{t-1}. \quad (11)$$

Following Deaton (1991) and Carroll (1997), I denote the sum of date t liquid wealth plus date t labor income by *cash-on-hand*. Period t cash-on-hand is equal to $X_t = LW_t + Y_t$ (in period 1 there is no initial level of housing and $LW_1 = 0$). At each date $t \leq T$, the investor needs to decide on the level of housing, consumption of other goods, whether to pay the fixed cost of equity market participation (if he has not done so before), and portfolio composition among liquid assets. I let FC_t take the value of 1 if the investor chooses to pay the fixed cost of equity market participation in period t , and zero otherwise, and denote by α_t the proportion of liquid assets invested in stocks over stocks plus bills, which by assumption is constrained to lie in the unit interval. The date t budget constraint is given by:

$$\begin{aligned} & S_t + B_t \\ &= \begin{cases} X_t - C_t - FC_t F - \delta P_t H_{t-1} + D_t, & \forall t \text{ s.t. No House Trade,} \\ X_t - C_t - FC_t F - \delta P_t H_{t-1} + D_t + (1 - \lambda) P_t H_{t-1} - P_t H_t, & \forall t \text{ s.t. House Trade.} \end{cases} \end{aligned} \quad (12)$$

Wealth at date $T+1$ is given by:

$$W_{T+1} = X_{T+1} - \delta H_T P_{T+1} + (1 - \lambda) H_T P_{T+1}. \quad (13)$$

The investor maximizes Equation (1) subject to Equations (2) through (13), plus the constraints that consumption must be nonnegative at all dates. The control variables for this problem are $\{C_t, H_t, D_t, \alpha_t, FC_t\}_{t=1}^T$, and the state variables are $\text{State}_t = \{t, H_{t-1}, X_t, \eta_t, \text{InvMove}_t, \text{IFC}_t\}_{t=1}^T$, where IFC_t is a state variable that takes the value of 1 if the investor is an equity market participant in period t and zero otherwise.

The Bellman equation for this problem is:

$$V_t(\text{State}_t) = \max_{C_t, H_t, D_t, FC_t, \alpha_t} \left[\frac{(C_t^{1-\theta} H_t^\theta)^{1-\gamma}}{1-\gamma} + \beta E_t V_{t+1}(\text{State}_{t+1}) \right] \quad \forall t \leq T,$$

where, given my assumptions, the level of housing at the beginning of date t is equal to the chosen level of housing at date $t-1$.

to pay down their debt, since if they are hit by a negative income shock they will be able to costlessly increase it [up to constraint (10)]. Campbell and Cocco (2003) study the choice between the standard fixed and adjustable rate mortgage contracts for a given house size.

1.6 Solution technique

The problem cannot be solved analytically. I use standard numerical techniques for solving it [Judd (1998)]. Given the finite nature of the problem, a solution exists and can be obtained by backward induction. I start by approximating the state-space and the variables over which the choices are made with equally spaced grids. The density functions for the random variables (namely, the innovations to the risky asset returns, house prices, and labor income process) were approximated using Gaussian quadrature methods to perform numerical integration [Tauchen and Hussey (1991)]. I use a three-state transition probability matrix to approximate the aggregate labor income process.

In period $T+1$, and for each admissible combination of the state variables, I obtain the utility associated with each level of terminal wealth. Since this is the terminal period, the utility function coincides with the value function. In every period t prior to $T+1$, I obtain the utility associated with the different choices of housing, other consumption, debt and portfolio choice among liquid assets. The date t value function is equal to current utility plus the expected discounted continuation value associated with the choices made, and given the value of the state variables. To compute this continuation value for points, which do not lie on the grid, I use cubic spline interpolation. The combinations of the choice variables ruled out by the constraints of the problem were attributed a very large (negative) utility so that they will never be optimal. I optimize over the different choices using grid search. I then iterate backward. Whenever, in the solution to the problem, the upper limit for the grids turned out to be binding, I increased it and solved the problem again.

2. Parameterization

2.1 Time parameters

Each period in the model is calibrated to correspond to five real life years.¹¹ I assume that the investor is born at age 25, retires at age 65, and dies at age 75. All parameters and labor income process were adjusted to take into account the five-year nature of each period.

2.2 Labor income

To parameterize the labor income process, I use data from the PSID for the years 1970 through 1992. The families that were part of the Survey of Economic Opportunities subsample were dropped to obtain a random sample of the U.S. population. The estimation was restricted to households with a male head.

¹¹ This is done for computational reasons, to keep the dimensionality of the problem low.




Figure 1
Five-year labor income profile

The data is from the PSID for the years 1970 through 1992. The families that were part of the Survey of Economic Opportunities were dropped from the sample. Labor income in each year is defined as total reported labor income plus unemployment compensation, workers compensation, social security, supplemental social security, other welfare, child support, and total transfers, all of which are for both the head of the household and, if present, his spouse. Labor income was deflated using the Consumer Price Index to obtain real variables. The figure plots the results for three education groups, according to the education of the head of the household.

In order to implicitly allow for (potentially) endogenous ways of self insurance against pure labor income risk, I use a broad definition of labor income. As in Cocco, Gomes, and Maenhout (2004), labor income in each year is defined as total reported labor income plus unemployment compensation, workers compensation, social security, supplemental social security, other welfare, child support, and total transfers (mainly help from relatives), all of which are for both the head of the household and, if present, his spouse. Labor income defined in this way was then deflated using the Consumer Price Index with 1992 as baseyear. Five-year labor income is for each household equal to the sum of discounted labor income over the relevant age group.¹² All age groups for which five observations were not available were dropped from the sample.

To estimate Equation (2), the function $f(t, Z_t)$ was assumed to be additively separable in t and Z_t , where the vector Z_t included age dummies, a family fixed effect, marital status, and household composition. As in Davis and Willen (2000), to control for education, the sample was split in three groups according to the level of education of the head of the household: no highschool degree, highschool degree, and college degree. Figure 1 shows the estimated age dummies for these education groups. These are the age profiles passed on to the simulation exercise.

¹² The rate used to discount labor income is 5%. This is the same rate that Heaton and Lucas (2000a) use.

The residuals obtained from the fixed-effects regressions of (log) labor income on $f(t, Z_{it})$ can be used to estimate σ_η^2 and σ_ω^2 . Define Y_t^* as:

$$\log(Y_{it}^*) \equiv \log(Y_{it}) - \hat{f}(t, Z_{it}). \quad (14)$$

Using Equation (3) to substitute in Equation (14) gives:

$$\log(Y_{it}^*) = \tilde{\eta}_t + \tilde{\omega}_{it}. \quad (15)$$

Averaging across all individuals gives:

$$\overline{\log(Y_{it}^*)} = \tilde{\eta}_t. \quad (16)$$

The variance of $\tilde{\eta}_t$ is obtained immediately as the variance of $\overline{\log(Y_t^*)}$. Subtracting this variance from the variance of $\tilde{\omega}_{it}$ for each education group gives the variance of $\tilde{\omega}_{it}$. The estimated standard deviations are shown in Table 1. This table also reports the estimated autoregressive coefficient in Equation (4). Labor income at retirement is set equal to the average of the labor income variable for the retirees in each education group in PSID data.

2.3 Housing

There are several parameters related to housing that I need to calibrate: the probability of an involuntary house trade (π), the minimum house size ($\bar{H}_{\min}$), and the parameters for the stochastic process for house prices.

I choose π so that the frequency of total (endogenous and exogenous) house transactions generated by the model matches the frequency of total

Table 1
Estimated parameters of the labor income and house price processes

Description	Parameter	Value
Autoregression parameter	ϕ	0.748
SD idiosyncratic inc. shocks		
No highschool degree	$\sigma_{\omega,1}$	0.136
Highschool degree	$\sigma_{\omega,2}$	0.131
College	$\sigma_{\omega,3}$	0.133
SD aggregate inc. shocks	σ_η	0.019
Real house price growth	b	0.016
SD house prices	$\sigma_{p'}$	0.062
Corr. house prices and agg. inc. shocks	$\rho_{\eta p'}$	0.553*
Corr. house prices and temp. inc. shocks	$\rho_{\omega p'}$	0.000

This table reports the estimated labor income and house price parameters. The data is from the the PSID for the years 1970 through 1992. The families that were part of the Survey of Economic Opportunities were dropped from the sample. Labor income in each year is defined as total reported labor income plus unemployment compensation, workers compensation, social security, supplemental social security, other welfare, child support, and total transfers, all of which are for both the head of the household and, if present, his spouse. Labor income and house prices were deflated using the Consumer Price Index to obtain real variables. House prices is an index of house prices for the families in PSID data.

*denotes significant at the 2% level.

house transactions in PSID data. In this data, and for the years 1970 through 1992, the proportion of households that were homeowners in two consecutive years, and that who reported that they had moved houses since the previous year is 5.44%. Since in my model each period is calibrated to correspond to five real life years, this implies a per period 24.4% moving probability. In the benchmark case, the transition matrix for the forced house sale is i.i.d., but in Section 4.6, I discuss how the results are affected if involuntary sales correlate with idiosyncratic labor income shocks.

Homeowners in the PSID are asked to assess the current (at the date of the interview) market value of their house.¹³ In order to parameterize minimum house size, I consider the house value reported by homeowners in the 1992 wave of the PSID, which is the most recent one available in final release form. Perhaps, surprisingly, there are homeowners who report fairly low house values: the first, fifth and tenth percentiles of reported house value distribution are U.S. \$2937, 11,380, and 22,026, respectively. Some of these reported house values probably are reporting errors. With this in mind, I set minimum house value in the model $\bar{H}_{\min}$ equal to U.S. \$20,000.

To estimate the parameters for the stochastic process for house prices, the self-assessed value of the house was deflated using the Consumer Price Index, with 1992 as the base year, to obtain real house prices. Define $p_{it} \equiv \log(P_{it})$, where P_{it} is the real price of house i at time t . Averaging across houses, I obtain for each year t an index of house prices:

$$p_t = \frac{\sum_{i=1}^{N_t} P_{it}}{N_t}, \quad t = 1970, \dots, 1992, \quad (17)$$

where N_t is the number of observations at time t . During this period, log real house prices increased an average 1.59% per year. Part of this increase is probably due to an improvement in the quality of housing, which cannot be accounted for using PSID data. Therefore, in the simulations, I decided to use a lower value for the average annual increase in house prices of 1%. Aggregate labor income is strongly positively correlated with cyclical fluctuations in house prices: the coefficient of correlation is as high as 0.553 and significantly different from zero.

The reader may be concerned that I construct an equally weighted house price index, instead of constructing an index using repeat sales of houses. The reason why I do so is that in PSID data only the current market value of the house is reported. When a household reports it has

¹³ A major concern with self-assessed values is that households, when asked about the current market value of their house, do not try to rationally assess this value. However, Skinner (1994) compared the self-assessed house values in the PSID to the objective measures of the Commerce Department and found that the two series are quite close in mapping housing price changes in the 1970s and 1980s.

moved houses since the previous interview, it reports the current market value of the new house, and not the value for which the old house was sold or for which the new house was purchased. This prevents me from constructing a house price index using repeat sales of houses.

It is of interest to compare the index I construct from PSID data to other house price indices. Figure 2 compares the Conventional Mortgage Home Price Index (CMHPI), constructed jointly by Freddie Mac and Freddie Mae, and the House Price Index (HPI), constructed by the Office of Federal Housing Enterprise Oversight, to the house price index I construct from PSID data. Both the CMHPI and the HPI indices are available since 1975 and are constructed using repeat sales of houses. Any time a house's value is observed twice over time (via either a sale or an appraisal), the change in the house price contributes to one observation of house price growth over that time period. The index is defined to be the statistically determined set of values that most closely fits such repeat observations [see Wang and Zorn (1997) for a description of the methodology employed]. In addition to the house price indices, Figure 2 plots the yearly percentage change in these indices.

Figure 2 shows that the index I construct from PSID data tracks other house price indices reasonably well. This is also reflected in the degree of correlation between the annual percentage changes in the PSID index and the annual percentage changes for the HPI and CMHPI indices: 0.922 and 0.871, respectively.




Figure 2
House price indices

This figure compares the house price index constructed using PSID data to the house price indices constructed using repeat sales data by the Freddie mac and the Office of Federal Housing Enterprise Overnight. The figure plots both the house price indices (1981 = 100) and the yearly percentage change in the indices.

2.4 Financial assets and credit markets

The riskfree rate used is 2% per year. To parameterize the mortgage rate premium, I subtract the three month Treasury bill rate to the 30-year fixed FHA nominal mortgage rate. The average annual premium for the period 1964–1997 is 3.01%. Part of this premium is the acquisition of an option on future inflation, a feature absent from the model. Accordingly, I decided to use a lower annual mortgage rate premium, of 2%, or an annual mortgage rate of 4%.

I use an annual mean return on risky assets of 10% and an SD of the log stock return equal to 0.1674 [Campbell, Lo and MacKinlay (1997)]. The one-time monetary fixed cost of equity market participation is set to \$1000, but I will consider other values for this parameter. The correlation between innovations to aggregate labor income and innovations to stock returns is slightly positive in the data, and equal to 0.047, although not statistically different from zero. Therefore, in the baseline case, I set it equal to zero.

2.5 Other parameters

In the baseline case, γ is equal to 5, below the upper bound of 10 considered plausible by Mehra and Prescott (1985). The parameter θ measures how much the investor values housing consumption relative to other goods consumption. Other preference parameters include the discount factor β . I parameterize these to at least roughly match the mean levels of housing and mortgage relative to financial assets observed in the data. Accordingly, the parameters chosen in the baseline case were β equal to 0.96 on an annual basis and θ equal to 0.10. I will do comparative statics with respect to these parameters. It may be reasonable to assume that θ varies over life, depending on the number and age of the children in the household. Introducing a time varying θ would be easy, but I abstract from these additional effects.

With respect to the transaction costs of changing houses, Smith, Rosen, and Fallis (1988) estimate for home ownership, the monetary component of these costs to be approximately 8–10% of the unit being exchanged. This estimate comprises transactions costs associated with search, legal costs, costs of readjusting home furnishings to a new house, and a psychic cost from disruption. The latter type of cost may vary significantly over life. As with θ , it would be fairly easy to consider age varying λ , but I abstract from this. I set λ equal to 8%. Reasonable values for the down payment are between 10 and 20% of the value of house. I use d equal to 15%. Leigh (1980) estimates the annual depreciation rate of housing units in the United States to be between 0.0036 and 0.0136. I use δ equal to 0.01 on an annual basis. Table 2 summarizes the parameters used in the baseline case.

Table 2
Baseline parameters (annual)

Description	Parameter	Value
Risk aversion	γ	5
Discount factor	β	0.96
Preference for housing	θ	0.10
Down payment	d	0.15
Depreciation rate	δ	0.01
Involuntary move	π	0.03
Transaction cost	λ	0.08
Riskless rate	$\bar{R}_F - 1$	0.02
Mortgage rate	$\bar{R}_D - 1$	0.04
Mean stock return	$\exp(\mu + \sigma_\eta^2/2) - 1$	0.10
Std of log stock return	σ_η	0.1674
Fixed cost	F	\$1000

This table reports the parameters used in the baseline case.

3. Results

3.1 Simulated data

To study the determinants of portfolio composition in the presence of housing, I simulate the behavior of investors in the model. I do so in two different ways. First, I generate different realizations for the temporary and the aggregate labor income shocks so that simulated households differ in the history of *both* of these shocks. The aim is to obtain the unconditional portfolio allocations generated by the model that is not conditional on one realization of the history of the aggregate labor income shock. Tables 4, 5, and 7 report these unconditional portfolio allocations.

However, these unconditional portfolio allocations may not be easily compared to the portfolio allocations that we observe in the data, since the latter are conditional on *one* realization of the history of aggregate labor income and house prices. Therefore, when I study the determinants of portfolio allocation using regression analysis (Section 4.4 and Table 6), I generate simulated data conditional on *one* realization of aggregate labor income and house prices. In this case, households of the same age differ in their expected earnings because of differences in educational attainment, and they also differ in the history of temporary labor income shocks, but they share the same realization of aggregate labor income and house price shocks. This realization is the one that most closely resembles PSID data in the years 1970 through 1989. For the years prior to 1970, which is relevant for simulated households older than 40 years of age in 1989, I set the realization of the aggregate shocks equal to its unconditional mean.

To make comparisons with PSID data meaningful, I group households in the 1989 cross-section of the PSID by age ranges and education levels and compute the relative weights of each age-education cohort. These

Table 3
Proportion of households by age and education

Age group	No highschool	Highschool	College	Total
25–29	0.01	0.06	0.02	0.09
30–34	0.02	0.08	0.04	0.14
35–39	0.01	0.10	0.06	0.17
40–44	0.02	0.08	0.07	0.16
45–49	0.01	0.05	0.03	0.09
50–54	0.02	0.04	0.02	0.07
55–59	0.02	0.04	0.02	0.08
60–64	0.02	0.03	0.02	0.08
65–69	0.02	0.04	0.01	0.07
70–74	0.02	0.03	0.01	0.06
Total	0.17	0.54	0.29	1.00

This table shows the proportion of households in each age and education group in the 1989 cross-section of the PSID. The families that were part of the Survey of Economic Opportunities were dropped from the sample.

weights are shown in Table 3.¹⁴ I then simulate a number of realizations that is in the same proportion in the simulated sample as in the data.

All components other than age and education in the vector Z_{it} were set equal to zero, so that the labor income profiles that I plot in Figure 1 are for a single investor without children. Of course, one could generate investors that differ in terms of their family composition and use the corresponding labor income profile. However, doing so is not trivial because when household composition is different from a single individual, the definition of consumption will have to be adjusted. To simplify, I abstract from this.

One important calibration issue that arose when solving the model is that when I increase the probability of a forced house sale, the frequency of voluntary house sales decreases. This makes sense since households are less likely to move voluntarily in a given period when there is a higher probability that they will be forced to move and pay the transaction cost of changing houses in the next period. For this reason, and in order to match the frequency of total house trades observed in the data, I had to solve the model several times, for different values for the probability of a forced house sale (π). When π is set equal to 3.2%, the frequency of total annual house sales is 5.5%, which roughly matches the value observed in the data.

One could also ask that the model matches the frequency of house sales conditional on age. With this in mind, I have looked at the frequency of house sales by age in PSID data and compared it to the frequency of house sales implied by the model. I have considered two age groups: households

¹⁴ I use the 1989 cross-section of the PSID since this is the one to which I compare the results of the model in Section 5.

whose head is less than 50 years of age, and households whose head is over 50 years of age. The annual probabilities of a house sale in PSID for these two age groups are 7.3 and 2.9%, respectively. The model implied that the annual probabilities of a house sale for the same age groups are 6.6 and 3.8%, respectively. Interestingly, and even though the assumed annual probability of a forced house sale is equal to 3.2% for both age groups, the model implied that the probability of a house sale is higher for the young than for the old. This happens for two reasons. First, young households are borrowing constrained, and the down payment constraint prevents them from buying the house size they desire in the first periods. Thus, the young are more likely to move voluntarily to upgrade the size of their house. Second, the bequest motive leads old investors not to wish to downgrade the size of the house they own; hence voluntary house sales are lower for the old.

I first examine the cross-sectional variation in the portfolio composition by net-worth and age generated by the model, and then do regression analysis. Following Heaton and Lucas (2000a), I consider several definitions of net-worth: (1) “liquid net-worth” is the sum of stocks and bills minus debt, (2) “financial net-worth,” which is liquid net-worth plus house value, and (3) “total net-worth,” which is financial net-worth plus capitalized labor and pension income. The terms “liquid assets,” “financial assets,” and “total assets” refer to the same classifications but without the subtraction of debt.

3.2 Portfolio composition by net-worth

A well-documented feature of the data on portfolio composition is that many households, particularly poorer ones, do not own stocks [Mankiw and Zeldes (1991)]. How does the investment in housing affect households’ willingness to participate in equity markets and the patterns of portfolio composition?

The first two columns of Table 4 show the mean share of stocks and bills in liquid assets, by financial net-worth, predicted by the model. I consider two financial net-worth groups: less than and greater or equal to U.S. \$100,000. When calculating the numbers in Table 4, one issue that has to be dealt with is that some investors who are not stock market participants do not hold any liquid assets. Strictly speaking, when liquid assets are zero, α is not defined, but in order to capture the fact that these investors are not stock market participants, I set α equal to zero.¹⁵

In the model, as in the data, low financial net-worth households do not participate in equity markets. Poorer households are liquidity constrained and, given their limited liquid wealth, prefer not to pay the fixed cost of

¹⁵ This is discussed further in Section 6.

Table 4
Portfolio shares by financial net-worth predicted by the model

Asset	Liquid assets		Financial assets		Total assets	
	<100k	≥100k	<100k	≥100k	<100k	≥100k
Stocks	0.252	0.952	0.049	0.460	0.017	0.213
Bills	0.748	0.048	0.029	0.013	0.008	0.006
Liquid assets	1.000	1.000	0.078	0.473	0.025	0.219
Real estate			0.922	0.527	0.248	0.251
Financial assets			1.000	1.000	0.273	0.470
Human capital					0.727	0.530
Total assets					1.000	1.000
Debt			0.509	0.034	0.141	0.018
Stock mkt part.	0.253	0.974				

This table reports mean portfolio shares of various assets relative to liquid assets, financial assets, and total assets. Liquid assets are the sum of stocks and Treasury bills. When liquid assets are zero, I set the ratio of stocks to liquid assets to zero to capture the fact that these investors are not stock market participants. Financial assets are liquid assets plus house value. Total assets are financial assets plus human capital. Debt is reported relative to financial assets and total assets. Stock market participation is the proportion of investors who participate in equity markets. Data are from simulating the model in Section 2 with the parameters shown in Table 2. Investors are categorized by financial net-worth. Financial net-worth is defined as the sum of stocks, bills, and house value less debt.

equity market participation. This results in a portfolio heavily tilted toward real estate: the mean real estate share is roughly 92% of financial assets for low financial net-worth households, which is much larger than the 53% for the high financial net-worth group. Poorer investors also tend to hold a highly leveraged portfolio, with a mean ratio of debt to financial assets as high as 51%.

3.3 Portfolio composition over the life-cycle

The theoretical literature on portfolio composition in presence of non-tradable income has found that when labor income shocks are uncorrelated with stock returns labor income resembles more closely Treasury bills than stocks [Jagannathan and Kocherlakota (1996), Heaton and Lucas (1997)]. This has implications for the composition of an investor's portfolio over the life-cycle: as investors age, implicit holdings of Treasury bills under the form of future labor income become less important, and investors make up for this decrease by shifting portfolio allocation toward riskless bills [Viceira (2001), Cocco, Gomes, and Maenhout (2004)]. Thus, the theoretical literature predicts that the portfolio share invested in stocks is decreasing over life. However, the empirical literature has found that the portfolio share invested in stocks is actually *increasing* over life, with some mixed evidence that points to slight decrease late in life. Can the investment in housing be the reason why the model without housing predicts a life-cycle pattern of stockholdings that is the opposite of what we observe in the data?

To address this question, the first four columns of Table 5 show the evolution over the life-cycle of the shares of stocks and bills relative to liquid assets, predicted by the model. The model with housing predicts an increasing life-cycle share of stock investments. Early in life, investment in housing keeps liquid assets low, and investors choose not to pay the fixed cost required for participating in the equity market. For investors in the lower age group, liquid assets are only 3% of financial assets. It is only later in life when liquid assets become sufficiently large, that stock market participation becomes more widespread.

Late in life, the presence of housing also prevents a decline in the share of stocks in liquid assets. In the model with housing, as investors age, liquid assets are less important relative to other asset holdings (human capital and housing) for future consumption. Old investors are more willing to accept risk in their portfolio of liquid assets since future consumption is less correlated with the return on the liquid assets portfolio.

Table 5 also shows that stock holdings are much less important when measured relative to financial assets rather than liquid assets. The reason is the importance of real estate, which varies between 97% for the youngest investors and 72% for investors in the 55–65 age group.

These magnitudes change again considerably when we consider total assets. Considering human capital as an asset provides the most complete view of investors' wealth. Human capital is computed in the simplest way possible: at each age, it is equal to the expected value of future labor

Table 5
Portfolio shares by age predicted by the model

Asset	Liquid assets				Financial assets				Total assets			
	<35	35–50	50–65	≥65	<35	35–50	50–65	≥65	<35	35–50	50–65	≥65
Stocks	0.019	0.213	0.613	0.716	0.005	0.060	0.235	0.096	0.001	0.013	0.099	0.072
Bills	0.981	0.787	0.387	0.284	0.025	0.024	0.041	0.011	0.003	0.005	0.016	0.008
Liquid assets	1.000	1.000	1.000	1.000	0.030	0.084	0.276	0.107	0.004	0.018	0.115	0.080
Real estate					0.970	0.916	0.724	0.893	0.126	0.148	0.293	0.715
Financial assets					1.000	1.000	1.000	1.000	0.130	0.166	0.408	0.795
Human capital									0.870	0.834	0.592	0.205
Total assets									1.000	1.000	1.000	1.000
Debt					0.683	0.428	0.265	0.512	0.088	0.067	0.117	0.417
Stock mkt part.	0.020	0.212	0.620	0.729								

This table reports mean portfolio shares of various assets relative to liquid assets, financial assets, and total assets. Liquid assets is the sum of stocks and Treasury bills. When liquid assets are zero, I set the ratio of stocks to liquid assets to zero to capture the fact that these investors are not stock market participants. Financial assets are liquid assets plus house value. Total assets are financial assets plus human capital. Debt is reported relative to financial assets and total assets. Stock market participation is the proportion of investors who participate in equity markets. Data are from simulating the model in Section 2, with the parameters shown in Table 2. Investors are categorized by age.

income discounted at the annual rate of 5%.¹⁶ The last four columns of Table 5 show that human capital is an important component of wealth at all ages, but particularly so for young investors: the share of human capital in total assets is as high as 87% for investors in the less than 35 age group.

When we consider total assets, there are also some striking changes in the life-cycle patterns of asset allocation: when measured relative to financial assets, the importance of real estate is roughly decreasing over life, whereas when measured relative to total assets it increases in importance throughout life. These changes in patterns arise naturally in my model due to the declining importance of capitalized labor income as investors age.

3.4 Determinants of portfolio composition

To more systematically summarize the correlations between stock holdings and other variables predicted by the model, I run the following Tobit regressions, where the dependent variable in the first equation is stocks relative to liquid assets (LA), financial assets (FA), and total assets (TA), and in the second equation is the dollar amount held in stocks:

$$\begin{aligned} \left(\frac{\text{Stocks}}{j} \right)_i &= a_1 + a_2 \text{INC}_i + a_3 \text{FNW}_i + a_4 \text{AGE}_i + a_5 \text{REFNW}_i \\ &\quad + a_6 \text{MORTFNW}_i + \delta_i \\ \text{Stocks}_i &= a_1 + a_2 \text{INC}_i + a_3 \text{FNW}_i + a_4 \text{AGE}_i + a_5 \text{RE}_i \\ &\quad + a_6 \text{MORT}_i + \delta_i, \end{aligned}$$

where $j = \text{LA, FA, TA}$. The independent variables are: income (INC), financial net-worth (FNW), age, real estate over financial net-worth (REFNW), and mortgage debt over financial net-worth (MORTFNW). MORT and RE are the corresponding dollar variables. The coefficients a_i are the regression coefficients and δ_i is the residual.¹⁷

Table 6 shows the correlations predicted by the model. In the model, labor income, although risky, resembles more closely bills than stocks. Therefore, higher levels of labor income induce a shift in portfolio allocation toward stocks. Hence the positive correlation between income and stock holdings.¹⁸

To understand the predicted correlation between the share of stocks held in liquid assets and financial net-worth, it is helpful to have in mind

¹⁶ This is the rate that Heaton and Lucas (2000a) use to compute the present discounted value of future labor income.

¹⁷ These are regressions similar to the ones run by Heaton and Lucas (2000a).

¹⁸ It is future and not current labor income that constitutes the implicit holdings of an asset. But the two tend to be positively correlated.

Table 6
Determinants of portfolio choice predicted by the model

	Stock relative to liquid assets	Stock relative to financial assets	Stock relative to total assets	Stocks
Intercept	-8.405 (-1.62)	0.423 (13.43)	0.085 (7.77)	-94223.46 (-19.57)
Total income	1.59e-05 (1.09)	3.73e-07 (5.78)	-1.67e-07 (-5.74)	0.209 (10.23)
Financial net-worth	2.92e-04 (4.09)	1.58e-06 (11.40)	1.35e-06 (22.01)	
Age	4.608 (7.05)	0.021 (15.91)	0.014 (17.78)	12110.85 (28.47)
Relative real estate	-48.834 (-6.14)	-0.762 (-32.56)	-0.266 (-27.11)	
Relative mortgage	48.174 (6.05)	0.725 (32.22)	0.264 (27.87)	
Real estate				0.245 (4.63)
Mortgage				-0.957 (-23.68)
Wald chi2(5)	52.58	12488	9652	1453

This table reports the results of cross-sectional Tobit regressions of several measures of stock holdings on several independent variables. The dependent variable is stock holdings relative to several measures of assets, and the dollar amount held in stocks. Data are from simulating the model in Section 2 with the parameters shown in Table 2. Total income is current labor income. Financial net worth is defined as the sum of stocks, bills, and house value less debt. Relative real estate is the value of the house relative to financial net-worth. Relative mortgage is the value of debt relative to financial net-worth. Real estate and mortgage are the corresponding dollar variables. Robust *T*-statistics are shown in parenthesis.

the shape of the underlying portfolio rule. For investors who have paid the fixed cost of equity market participation, the portfolio rule, conditional on given levels of housing and debt and at a given age, is decreasing in financial net-worth. Investors hold bills under the form of future labor income. As financial net-worth increases, these implicit holdings of Treasury bills become less important relative to financial wealth, and investors make up for this decrease by shifting portfolio composition toward bills.

For investors who have not yet paid the fixed cost of equity market participation, the portfolio rule is not monotonic. For low levels of financial net-worth, investors decide not to participate in equity markets. It is only when financial net-worth (and liquid assets) are sufficiently large that they pay the fixed cost, and at this point, stock holdings increase with financial net-worth. After this level, the decreasing portfolio rule described above obtains. Table 6 shows that for stocks relative to liquid assets, the increasing part of the rule tends to dominate so that the predicted correlation between stockholdings and financial net-worth is positive.

The effects of nonparticipation in equity markets are picked up by the real estate variable. Investors who have high levels of real estate relative to financial net-worth do not participate in equity markets. Hence the

negative correlation between these variables. Interestingly, mortgage debt tends to be positively correlated with stock holdings. The reason is that investors who are better educated and have higher expected future labor income borrow more, but at the same time, labor income resembles more closely bills than stocks, which induces a shift in the portfolio allocation toward stocks.

The coefficient on age is positive, which reflects the life-cycle patterns shown in Table 5. Table 6 shows that some of the predicted correlations are sensitive to the measure of stock holdings used. In particular, the coefficient on income flips from positive to negative when we measure stocks relative to total assets.

In Section 5, I will examine the cross-sectional variation in the composition of wealth by age and net-worth, and run regressions similar to those in Table 6 on PSID portfolio composition data to study the extent to which the effects at work in the model are also present in the data. This comparison will provide evidence on the strengths and limitations of the model.

3.5 The effects of house price risk

The housing investment keeps liquid assets low and poorer investors from participating in equity markets. But what are the effects of *house price risk* on asset allocation? In particular, does house price risk crowd out stock-holdings? Table 7 compares the portfolio shares for the baseline case to the case when the variance of house price shocks is set to zero (the no house price risk scenario). One may reasonably expect that house price risk has an impact on asset allocation that depends on financial net-worth. Therefore, Table 7 reports portfolio shares conditional on financial net-worth.

Table 7
The effects of house price risk

Asset	Baseline		No house price risk	
	<100k	≥100k	<100k	≥100k
Stocks	0.049	0.460	0.056	0.489
Bills	0.029	0.013	0.022	0.002
Liquid assets	0.078	0.473	0.078	0.491
Real estate	0.922	0.527	0.922	0.509
Financial assets	1.000	1.000	1.000	1.000
Debt	0.508	0.034	0.473	0.013
Stock mkt part.	0.253	0.974	0.311	1.000

This table reports mean portfolio shares of various assets relative to financial assets. Liquid assets are the sum of stocks and Treasury bills. Financial assets are liquid assets plus house value. Debt is reported relative to financial assets. Stock market participation is the proportion of investors who participate in equity markets. Data are from simulating the model in Section 2 with the parameters shown in Table 2, with the standard deviation of house price shocks set to the baseline case and to zero.

Table 7 shows that house price risk crowds out stock holdings, both for high and low financial net-worth investors. As expected, this crowding out effect is larger for lower financial net-worth investors: in the presence of house price risk, the portfolio share of stocks relative to financial assets is 13% lower for households with less than \$100,000, and only 6% lower for wealthier households. The relative importance of stock holdings in the liquid assets portfolio is also lower in the presence of house price risk, both for high and low financial net-worth investors. Finally, Table 7 shows that house price risk also leads to lower stock market participation with a larger quantitative effect for low financial net-worth households.

One perhaps surprising finding is that wealthy investors on average invest less in real estate when there is no house price risk compared to when there is house price risk. This is due to a sample selection issue. When there is house price risk, the sample of wealthy investors tends to be composed of those investors who face high current house prices (i.e., higher than when there is no house price risk), and who have bought their first house or upgraded the size of their house when house prices were relatively low (i.e., lower than when there is no house price risk). Thus, and since housing is such an important component of wealth, those investors who are fortunate to step into the housing ladder or to trade their houses for larger ones at low housing prices, are ex-post more likely to have higher financial net-worth and to invest a higher fraction of their wealth in real estate.¹⁹

3.6 Comparative statics

In this section, I discuss the effects of certain parameters of my model, including the correlation between temporary labor income shocks and house price risk, which I set equal to zero in the baseline case. The effects of this parameter are of particular interest since they may help us understand the effects of the assumed perfect correlation between aggregate labor income risk and house prices.²⁰

Positive correlation between the housing asset and labor income shocks makes the housing asset riskier since it is not as good a hedge against labor income risk. Therefore, households shift their financial portfolio toward liquid financial assets and away from real estate. The larger liquid financial assets makes it more worthwhile for households to pay the monetary fixed cost of equity participation, so that, perhaps surprisingly, stock

¹⁹ Since when there is house price risk wealthy investors have, on average, bought the house at a lower price, they have, on average, bought larger houses. Furthermore, since when there is house price risk wealthy investors tend to be those that face higher current house prices, on average they invest more in real estate than when there is no house price risk.

²⁰ I would like to thank an anonymous referee for suggesting to me that explore this issue.

market participation is larger in the presence of positive correlation. However, since housing enters directly into the utility function of agents, the crowding out of real estate is not very large: when I set the correlation between labor income and house price shocks equal to 0.30, the (unconditional) share of financial assets in real estate decreases by 1 percentage point, from 88 to 87%. The proportion of households that are stock market participants increases from 33 to 34.3%, and the shares of stocks and bonds in the financial portfolio increase from 9.4 and 2.6, to 9.9 and 3.1%, respectively. Since the percentage increase in bond holdings is larger than that in stockholdings, there is a tilt in the liquid assets portfolio toward the riskless asset.

Another parameter that plays an important role in my model is the fixed cost of equity market participation. Without it, every investor in the model would own stocks, which would be at odds with the well-documented feature of the data on portfolio composition that many households do not own stocks at all. It is obvious that fixed costs can generate limited stock market participation. The important question is the *magnitude* of the fixed cost needed to generate the levels of stock market participation observed empirically. A fixed cost of U.S. \$1000 goes a long way toward generating limited stock market participation: the proportion of investors who participate in equity markets drops from 47% for a fixed cost of \$500, to roughly 33% for a fixed cost of \$1000. Since the housing investment keeps liquid assets and the benefits of stock market participation small, a lower fixed cost is needed to generate the levels of participation observed empirically.

One could argue that liquid assets and participation rates are low early in life because of consumption smoothing, and not necessarily because of the housing investment. To evaluate the extent to which the housing investment contributes to the low participation rates, I have solved the model ignoring housing, but keeping all other parameters at their benchmark values, including the \$1000 fixed cost of equity market participation. In this case, the proportion of households who participate in equity markets is equal to 76%, which is much higher than the 33% predicted by the model with housing.

The minimum house size restriction also plays a role in preventing households from participating in equity markets and in lowering stockholdings. The larger is $H_{\min}$, the lower is the stock market participation. For a minimum house size of \$40,000, which is double the benchmark value, the proportion of households who participate in equity markets drops from 33 (the benchmark value) to 21%, and the share of stocks in financial assets drops from 9.3 to 6.3%.

Transaction costs of adjusting the level of housing are substantial, and since investors trade owner-occupied houses for the purpose of consumption smoothing as well as portfolio rebalancing, they incur these costs

frequently.²¹ The impact of higher transaction costs is similar to that of an increased probability of a forced move in that both lead to a reduction in the frequency of endogenous house trades.

The effects of transaction costs and of the probability of a forced move on stock market participation are not as straightforward as the effects of the minimum house size. When the probability of a forced move is higher, so that investors are forced to pay the transaction costs of housing adjustment more often, they react to this added risk by buying smaller houses, and shifting portfolio composition toward liquid financial assets (both stocks and bonds). The higher liquid financial assets make households more willing to pay the fixed cost of equity market participation.

The effects of the probability of a forced house sale are particularly large when they are negatively correlated with idiosyncratic labor income shocks: the negative correlation makes it more likely that households have to pay the (large) transaction costs associated with a house sale in those states where income is also low. When the correlation is equal to -1 , the mean share of real estate in the financial portfolio decreases by as much as 30%. At the same time, and because of the higher liquid financial assets, stock market participation increases to roughly 58%.

Preference parameters also have effects on portfolio composition and stock holdings. The nature of these effects depends on the parameter considered. For a lower discount factor, investors save less and are less willing to pay the fixed cost of equity market participation. The effects of the discount factor are particularly large in my model since liquid assets are kept low by the housing investment. Lower stock market participation can also be obtained by increasing the importance of housing as a consumption good (i.e., by increasing θ). Households invest more in real estate and hold more leveraged portfolios.

4. Comparison of the Model with the Data

To study the extent to which the effects at work in the model are also present in the data, I compare portfolio composition predicted by the model to that observed in PSID data.²² I use the 1989 wave that contains asset information. Throughout, I restrict the sample to those households

²¹ The effects of transaction costs of portfolio adjustment have been studied by Balduzzi and Lynch (1999), Constantinides (1986), Davis and Norman (1990), Heaton and Lucas (1997), among others. The literature has focused on transaction costs of adjusting stock holdings. Transaction costs tend to have larger effects on portfolio allocations when investors trade for the purpose of consumption smoothing as well as portfolio rebalancing.

²² Recent empirical studies on portfolio composition include Bertaut and Haliassos (1997), Goetzmann (1993) Guiso, Jappelli, and Terlizzese (1996), Heaton and Lucas (2000a), Poterba and Samwick (1997), or see Heaton and Lucas (2000b) for an excellent survey of this literature, and Ameriks and Zeldes (2001) who study the empirical relationship between age and the fraction of wealth invested in the stock market.

who own a house,²³ and do not belong to the Survey of Economic Opportunities. I first describe the variables used, and then present and discuss the results.

4.1 Description of the data

When considering data on portfolio composition, there exist asset categories that are not present in the model. Broadly these include bonds, vehicles, real estate other than the main home, and business assets.²⁴ This requires that we restate our definitions of net-worth: (1) “liquid net-worth” is the sum of stocks, bonds, and cash²⁵ minus all forms of debt, (2) “financial net-worth,” which is liquid net-worth plus house value, vehicles, other real estate, and the value of family owned business or farm, and (3) “total net-worth,” which is financial net-worth plus capitalized labor and pension income. As before, the terms “liquid assets,” “financial assets,” and “total assets” refer to the same classifications but without the subtraction of debt.²⁶

An important asset for most households is human capital. I compute the value of human capital in the simplest way possible. For those households whose head is younger than 65 years of age, broadly defined labor income is assumed to remain constant until age 65, at which age it decreases in the average proportion of the decrease observed in the data. For households whose head is over 65-years of age, labor income was assumed to remain constant until age 75. Human capital is equal to the present discounted value of future labor income, discounted at the annual rate of 5%.²⁷

4.2 Portfolio composition by financial net-worth and over the life-cycle

I first consider the cross-sectional variation in the composition of wealth by financial net-worth and age, and compare it to that predicted by the model. Table 8 shows that the shares of liquid and financial assets held in stocks increase with financial net-worth. The levels in the data are smaller than those predicted by the model, both for low and high financial net-worth investors, although for high financial net-worth households the percentage difference is larger. When considering stock holdings over

²³ This may introduce biases as households who own a house tend to be older and wealthier. Households whose heads were older than 75 years of age were dropped from the sample.

²⁴ The PSID data does not contain information on the value of pensions and retirement plans. Retirement wealth was also treated in a stylized way in the model.

²⁵ I consider cash to include money in checking or savings accounts, money market funds, Treasury bills, and certificates of deposit. These assets are riskless in nominal terms whereas the riskless asset in the model is riskless in real terms. I use these definitions as an approximation.

²⁶ All households who refused to answer or did not know the answer to the amounts invested in these assets and those for whom financial net-worth is negative were dropped from the sample.

²⁷ A measure of human capital was computed in a similar manner by Heaton and Lucas (2000a).

Table 8
Portfolio shares by financial net-worth in PSID data

Asset	Liquid assets		Financial assets		Total assets	
	<100k	≥100k	<100k	≥100k	<100k	≥100k
Stocks	0.102	0.242	0.015	0.073	0.002	0.030
Cash	0.749	0.599	0.058	0.127	0.011	0.050
Bonds	0.149	0.159	0.022	0.040	0.003	0.015
Liquid assets	1.000	1.000	0.095	0.240	0.016	0.095
Real estate			0.762	0.488	0.119	0.169
Vehicles			0.109	0.064	0.014	0.021
Other real estate			0.023	0.104	0.005	0.049
Business			0.011	0.104	0.004	0.059
Financial assets			1.000	1.000	0.158	0.393
Human capital					0.842	0.607
Total assets					1.000	1.000
Debt			0.384	0.121	0.052	0.036
Stock mkt part.	0.179	0.543				

This table reports mean portfolio shares of various assets relative to liquid assets, financial assets, and total assets. Liquid assets are the sum of stocks, cash, and bonds. Cash includes money in checking or savings accounts, money market funds, Treasury bills, and certificates of deposit. Financial assets are liquid assets plus house value, vehicles, other real estate, and the value of family owned business. Total assets are financial assets plus human capital. Debt includes mortgage debt, and is reported relative to financial assets and total assets. Stock market participation is the proportion of investors who participate in equity markets. Data are from the 1989 wave of the PSID. Households are categorized by financial net-worth. Financial net-worth is equal to financial assets minus debt. The sample is restricted to households who own a house, and do not belong to the Survey of Economic Opportunities. All households with negative net-worth were dropped from the sample.

the life-cycle, Table 9 shows that the share of stocks is increasing over life, with a slight decrease after age 65 for stocks relative to liquid assets. Comparing Tables 5 and 9, we see that the model with housing is more successful at keeping the levels of stockholdings and equity market participation low among young households than among households older than 50 years of age.

Table 9 also shows that as in the model human, capital is an important component of wealth at all ages, but more so for younger households. As in the model, the declining importance of human capital over life explains why, when measured relative to total assets, the importance of real estate increases over life, whereas the reverse is true for real estate relative to financial assets. Leverage ratios decrease over life. Summarizing, the model is able to explain the cross-sectional patterns of portfolio composition by financial net-worth and age. The housing investment plays a crucial role in keeping stockholdings and equity market participation low, mainly among low net-worth and young households.

4.3 Determinants of portfolio composition

I run regressions similar to those run using synthetic data, and present the results in Table 10. I include business wealth in these regressions, since Heaton and Lucas (2000a) emphasize the importance of entrepreneurial

Table 9
Portfolio shares by age in PSID data

Asset	Liquid assets				Financial assets				Total assets			
	<35	35–50	50–65	≥65	<35	35–50	50–65	≥65	<35	35–50	50–65	≥65
Stocks	0.120	0.172	0.178	0.147	0.016	0.032	0.047	0.054	0.002	0.007	0.019	0.029
Cash	0.738	0.655	0.685	0.711	0.051	0.061	0.106	0.150	0.007	0.013	0.037	0.063
Bonds	0.142	0.173	0.137	0.143	0.024	0.031	0.031	0.027	0.004	0.007	0.010	0.011
Liquid assets	1.000	1.000	1.000	1.000	0.091	0.124	0.184	0.231	0.013	0.027	0.066	0.103
Real estate					0.732	0.686	0.609	0.597	0.092	0.122	0.164	0.199
Vehicles					0.114	0.096	0.082	0.069	0.013	0.015	0.022	0.022
Other real estate					0.024	0.046	0.073	0.073	0.004	0.015	0.033	0.038
Business					0.039	0.048	0.051	0.030	0.016	0.023	0.033	0.019
Financial assets					1.000	1.000	1.000	1.000	0.138	0.202	0.318	0.381
Human capital									0.862	0.798	0.682	0.618
Total assets									1.000	1.000	1.000	1.000
Debt									0.062	0.057	0.032	0.016
Stock mkt part.	0.257	0.344	0.324	0.268	0.503	0.346	0.145	0.060				

This table reports mean portfolio shares of various assets relative to liquid assets, financial assets, and total assets. Liquid assets are the sum of stocks, cash, and bonds. Cash includes money in checking or savings accounts, money market funds, Treasury bills, and certificates of deposit. Financial assets are liquid assets plus house value, vehicles, other real estate, and the value of family owned business. Total assets are financial assets plus human capital. Debt includes mortgage debt and other debt, and is reported relative to financial assets and total assets. Stock market participation is the proportion of investors who participate in equity markets. Data are from the 1989 wave of the PSID. Households are categorized by the age of the head of the household. The sample is restricted to households who own a house, and do not belong to the Survey of Economic Opportunities. All households with negative net-worth were dropped from the sample.

Table 10
Determinants of portfolio choice in PSID data

	Stock relative to liquid assets	Stock relative to financial assets	Stock relative to total assets	Stocks
Intercept	-0.257 (-2.04)	-0.119 (-2.58)	-0.071 (-3.56)	-34.818e+4 (-2.75)
Total income	3.43e-06 (2.47)	1.25e-06 (2.48)	3.44e-07 (1.64)	2.221 (1.55)
Financial net-worth	7.39e-08 (1.23)	3.13e-08 (1.40)	2.82e-08 (1.68)	
Age	0.002 (1.64)	0.002 (3.98)	0.001 (4.99)	1952.46 (2.83)
Relative real estate	-0.360 (-3.94)	-0.190 (-4.88)	-0.082 (-4.73)	
Relative mortgage	0.355 (4.01)	0.188 (4.92)	0.081 (4.77)	
Relative business	-0.373 (-4.13)	-0.196 (-5.33)	-0.081 (-4.83)	
Real estate				0.400 (4.87)
Mortgage				-0.066 (-0.24)
Business				-0.130 (-1.00)
Wald chi2(6)	35.89	63.28	62.44	422.23

This table reports the results of cross-sectional Tobit regressions of several measures of stock holdings on several independent variables. The dependent variable is stock holdings relative to several measures of assets, and the dollar amount held in stocks. Data are from the 1989 wave of the PSID. Total income is current labor income. Financial net-worth is financial assets minus debt. Age is the age of the head of the household. Relative real estate is the value of the house relative to financial net-worth. Relative mortgage is the value of mortgage debt relative to financial net-worth. Real estate, mortgage, and business are the corresponding dollar variables. The sample is restricted to households who own a house, and do not belong to the Survey of Economic Opportunities. All households with negative net-worth were dropped from the sample. Robust *T*-statistics are shown in parenthesis.

risk for portfolio choice. The estimated coefficients should be compared to those predicted by the model and shown in Table 6. Although the model is able to predict the signs of the estimate coefficients for stocks relative to liquid and financial assets and total stock holdings, there are differences for stocks relative to total assets, namely the sign of the estimated coefficient for income. These differences are discussed in the next section.

5. Discussion of Model and Results

The model is able to explain the patterns of cross-sectional variation in the composition of wealth observed in the data, and the effects of the housing investment on portfolio composition, but it has several limitations that need to be acknowledged and discussed.

In the parameterization section, I set $\exp(\mu + \sigma_v^2/2) > \bar{R}_D > \bar{R}_F$. Since debt levels can be costlessly renegotiated, it follows that no investor holds bills and debt simultaneously. Those investors who do not participate in equity markets prefer to pay their debt rather than hold bills, and stock

market participants prefer to pay their debt or invest in stocks rather than hold bills. Since, in practice, many debt holders (whether stock market participants or not) also hold bills, a natural question to ask is what is missing from the model that may explain this counterfactual implication. A natural candidate is the assumption that debt levels can be costlessly renegotiated. When it is costly to increase the level of outstanding debt, as, for example, when renegotiating a conventional mortgage contract, investors may wish to simultaneously hold a small amount of bills and debt.²⁸

Another way of generating simultaneous holdings of debt and bills is to assume that for transaction purposes investors have to hold cash equal to at least a given proportion of current nondurable consumption. This assumption may be motivated by the cash-in-advance models [see, e.g., Lucas (1982)]. It may also be motivated by the fact that for practical purposes, investors need to hold some cash to pay for the goods they consume. It amounts to the restriction that $B_t \geq \tau C_t$, where τ is the proportion of current nondurable consumption that must be held in bills. Introducing this restriction in the model is easy, but parameterizing τ would be difficult and to some extent arbitrary. When we let τ go to zero we obtain the results shown in Table 4.

A second important limitation of the analysis is that, although the model with housing lowers stock market participation and unconditional stock shares, it is much less successful at matching stock shares conditional on participation with predicted values much higher than those we observe in the data. This is particularly true for those households with financial net-worth of less than \$100,000 who participate in equity markets, for whom the portfolio shares of liquid, financial, and total assets invested in stocks predicted by the model are 100, 19.3, and 6.6%, respectively. These clearly are too high when compared to PSID data where the corresponding values are 46.4, 4.1, and 0.7%, respectively.²⁹

The assumption of costless debt adjustment plays a role in the model's inability to match stock shares conditional on stock market participation. If increasing debt levels were costly, smoothing income shocks by increasing debt levels would be suboptimal within a certain region of the state space, and stock market participants may wish to hold bonds in their liquid assets portfolio in order to hedge against income shocks. This would be particularly true if income shocks were positively correlated

²⁸ I have solved a modified version of the model, in which households have to pay a monetary cost of U.S. \$500 if they wish to increase debt above the current level. To be able to solve this model, which requires an additional expensive state variable (debt outstanding), I had to set the variance of house price shocks and aggregate labor income shocks, and the probability of a forced house sale equal to zero (this reduces the number of state variables by 2). I find that households who do not participate in the stock market wish to simultaneously hold debt and a small amount of bills equal to an average of \$730.

²⁹ The difficulty that the model has in generating low stockholdings conditional on participation and at the same time matching participation rates is not unique to my model [see Gomes and Michaelides (2004)].

with stock returns as shown by Heaton and Lucas (2000a) and Cocco, Gomes, and Maenhout (2004), and as estimated by Davis and Willen (2000) for some cohorts of the U.S. population.

For older and high financial wealth households, the model also predicts stock shares that are higher than in the data, even unconditionally. The comparison of Tables 4 and 8, and 5 and 9 help to explain why. From these tables we see that one important asset for these households, not considered in the model, is business wealth. More than 10% of the financial assets of high net-worth households are under the form of privately held businesses. Business wealth may help reduce the level of stock holdings because the income from businesses is more volatile and more correlated with stock returns than labor income [Polkovnichenko (1999), Heaton and Lucas (2000a)], and the portfolios of entrepreneurial households are very undiversified [Quadrini (1999), Hubbard and Gentry (2000)].

A further limitation of the analysis is the assumed stochastic process for house prices. I focused attention on cyclical fluctuations in house prices and the correlation between the latter, income shocks, and stock returns. However, the true stochastic process for house prices is likely to be more complex than the one I have assumed, involving higher-order autoregressive or moving average terms [Case and Shiller (1989), Poterba (1991)].³⁰

A final limitation of the analysis is taxes, which I ignored. It is intuitive that the tax benefits associated with mortgage interest payments may have important effects on asset allocation. For example, they may help magnify the effects of nontradable income on debt levels: investors who expect their future labor income to be higher, and be in a higher tax bracket, will borrow more.

6. Conclusion

In this article, I studied portfolio choice in the presence of housing. This is important since owner-occupied housing is the single most important asset in many investors' portfolios. The model provided answers to the questions raised in the introduction, which I now briefly discuss. Investment in housing has important implications for asset accumulation and portfolio choice among stocks and Treasury bills. Early in life, and at low levels of financial net-worth, it keeps liquid assets low and reduces the benefits of equity market participation. House price risk crowds out stockholdings, and this crowding out effect is larger for lower financial

³⁰ Of course, the difficulty in considering more general processes for house prices is that they lead to an increase in the number of state variables.

net-worth. The model also proposed an explanation as to why in the data leverage and stockholdings tend to be positively correlated. For investors with a more leveraged portfolio, capitalized labor income tends to be a more important component of wealth. Higher capitalized labor income induces a shift in portfolio composition toward stocks so that leverage and stock holdings tend to be positively correlated.

References

- Ameriks, J., and S. P. Zeldes, 2001, "How Do Household Portfolio Shares Vary with Age?" working paper, Columbia Business School.
- Balduzzi, P., and A. W. Lynch, 1999, "Transaction Costs and Predictability: Some Utility Cost Calculations," *Journal of Financial Economics*, 52, 47–78.
- Basak, S., and D. Cuoco, 1998, "An Equilibrium Model with Restricted Stock Market Participation," *The Review of Financial Studies*, 11, 309–341.
- Bertaut, C. C., and M. Haliassos, 1997, "Precautionary Portfolio Behavior from a Life-Cycle Perspective," *Journal of Economic Dynamics and Control*, 21, 1511–1542.
- Bodie, Z., R. C. Merton, and W. Samuelson, 1991, "Labor Supply Flexibility and Portfolio Choice in a Life Cycle Model," *Journal of Economic Dynamics and Control*, 16, 427–449.
- Campbell, J. Y., and J. F. Cocco, 2003, "Household Risk Management and Optimal Mortgage Choice," *Quarterly Journal of Economics*, 118, 1149–1194.
- Campbell, J. Y., A. W. Lo, and A. C. MacKinlay, 1997, *The Econometrics of Financial Markets*, Princeton University Press, Princeton, NJ.
- Carroll, C. D., 1997, "Buffer-Stock Saving and the Life-Cycle/Permanent Income Hypothesis," *Quarterly Journal of Economics*, 112, 1–55.
- Case, K. E., and R. J. Shiller, 1989, "The Efficiency of the Market for Single-Family Homes," *American Economic Review*, 79, 125–137.
- Cocco, J. F., F. J. Gomes, and P. J. Maenhout, 2004, "Consumption and Portfolio Choice over the Life-Cycle," working paper, London Business School; forthcoming in *The Review of Financial Studies*.
- Constantinides, G. M., 1986, "Capital Market Equilibrium with Transaction Costs," *Journal of Political Economy*, 94, 864–862.
- Cuoco, D., and H. Liu, 2000, "Optimal Consumption of a Divisible Durable Good," *Journal of Economic Dynamics and Control*, 24, 561–613.
- Davis, M., and A. Norman, 1990, "Portfolio Selection with Transaction Costs," *Mathematics of Operations Research*, 15, 676–713.
- Davis, S., and P. Willen, 2000, "Using Financial Assets to Hedge Labor Income Risks: Estimating the Benefits," working paper, University of Chicago.
- Deaton, A., 1991, "Savings and Liquidity Constraints," *Econometrica*, 59, 1221–1248.
- Faig, M., and P. Shum, 2002, "Portfolio Choice in the Presence of Personal Illiquid Projects," *Journal of Finance*, 57, 303–328.
- Flavin, M., and T. Yamashita, 2002, "Owner-Occupied Housing and the Composition of the Household Portfolio over the Life Cycle," *American Economic Review*, 92, 345–362.
- Goetzmann, W. N., 1993, "The Single Family Home in the Investment Portfolio," *Journal of Real Estate Finance and Economics*, 6, 201–222.

- Gomes, F., and A. Michaelides, 2004, "Optimal Life-Cycle Asset Allocation: Understanding the Empirical Evidence," working paper, London Business School; forthcoming in *Journal of Finance*.
- Grossman, S. J., and G. Laroque, 1991, "Asset Pricing and Optimal Portfolio Choice in the Presence of Illiquid Durable Consumption Goods," *Econometrica*, 58, 25–51.
- Guiso, L., T. Jappelli, and D. Terlizzese, 1996, "Income Risk, Borrowing Constraints and Portfolio Choice," *American Economic Review*, 86, 158–172.
- Heaton, J., and D. J. Lucas, 1997, "Market Frictions, Saving Behavior and Portfolio Choice," *Macroeconomic Dynamics*, 1, 76–101.
- Heaton, J., and D. J. Lucas, 2000a, "Portfolio Choice and Asset Prices: The Importance of Entrepreneurial Risk," *Journal of Finance*, 55, 1163–1198.
- Heaton, J., and D. J. Lucas, 2000b, "Portfolio Choice in the Presence of Background Risk," *Economic Journal*, 110, 1–26.
- Hu, X., 2002, "Portfolio Choice for Home Owners," working paper, Northwestern University.
- Hubbard, G., and W. Gentry, 2000, "Entrepreneurship and Household Saving," NBER Working Paper, 7894.
- Hubbard, G., J. S. Skinner, and S. Zeldes, 1994, "The Importance of Precautionary Motives for Explaining Individual and Aggregate Saving," in Allan H. Meltzer and Charles I. Plosser (eds.), *Carnegie-Rochester Conference Series on Public Policy*, 40, 59–125.
- Jagannathan, R., and N. R. Kocherlakota, 1996, "Why Should Older People Invest Less in Stocks than Younger People?" Federal Reserve Bank of Minneapolis, *Quarterly Review*, 20, 11–23.
- Judd, K. L., 1998, *Numerical Methods in Economics*, The MIT Press, Cambridge, Massachusetts.
- Leigh, W. A., 1980, "Economic Depreciation of the Residential Housing Stock of the United States, 1950–1970," *Review of Economics and Statistics*, 62, 200–206.
- Lucas, R., 1982, "Interest Rates and Currency Prices in a Two-Country World," *Journal of Monetary Economics*, 10, 335–359.
- Luttmer, E., 1999, "What Level of Fixed Costs Can Reconcile Consumption and Stock Returns?" *Journal of Political Economy*, 107, 969–997.
- Mankiw, N. G., and S. Zeldes, 1991, "The Consumption of Stockholders and Non-Stockholders," *Journal of Financial Economics*, 29, 97–112.
- Mehra, R., and E. Prescott, 1985, "The Equity Premium Puzzle," *Journal of Monetary Economics*, 15, 145–161.
- Merton, R., 1971, "Optimum Consumption and Portfolio Rules in a Continuous-Time Model," *Journal of Economic Theory*, 3, 373–413.
- Polkovnichenko, V., 1999, "Heterogeneity and Proprietary Income Risk: Implications for Stock Market Participation and Asset Prices," working paper, Northwestern University.
- Polkovnichenko, V., 2000, "Limited Stock Market Participation and the Equity Premium," working paper, University of Minnesota.
- Poterba, J. M., 1991, "House Price Dynamics: The Role of Tax Policy and Demography," *Brookings Papers on Economic Activity*, 2, 143–183.
- Poterba, J. M., and A. A. Samwick, 1997, "Household Portfolio Allocation over the Life Cycle," NBER Working Paper, 6185.
- Quadrini, V., 1999, "Entrepreneurship, Saving, and Social Mobility," *Review of Economic Dynamics*, 3, 1–40.

- Samuelson, P. A., 1969, "Lifetime Portfolio Selection by Dynamic Stochastic Programming," *Review of Economics and Statistics*, 51, 239–246.
- Skinner, J., 1994, "Housing and Saving in the United States," in Y. Noguchi and J. M. Poterba (eds.), *Housing Markets in the United States and Japan*, University of Chicago Press, Chicago, pp. 191–213.
- Smith, L. B., K. T. Rosen, and G. Fallis, 1988, "Recent Developments in Economic Models of Housing Markets," *Journal of Economic Literature*, XXVI, 29–64.
- Tauchen, G., and R. Hussey, 1991, "Quadrature-Based Methods for Obtaining Approximate Solutions to Nonlinear Asset Pricing Models," *Econometrica*, 59, 371–396.
- Viceira, L. M., 2001, "Optimal Portfolio Choice for Long-Horizon Investors with Nontradable Labor Income," *Journal of Finance*, 56, 433–470.
- Vissing-Jorgensen, A., 2002, "Towards an Explanation of Household Portfolio Choice Heterogeneity: Nonfinancial Income and Participation Cost Structures," working paper, Northwestern University.
- Wang, F. T., and P. Zorn, 1997, "Estimating House Price Growth with Repeat Sales Data: What's the Aim of the Game?" *Journal of Housing Economics*, 6, 93–118.
- Yao, R., and H. Zhang, 2004, "Optimal Consumption and Portfolio Choices with Risky Housing and Borrowing Constraints," working paper, University of North Carolina; forthcoming in *The Review of Financial Studies*.

Copyright of Review of Financial Studies is the property of Society for Financial Studies and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.