

Owner-Occupied Housing and the Composition of the Household Portfolio

By MARJORIE FLAVIN AND TAKASHI YAMASHITA*

For most homeowners, the house is the single most important consumption good appearing as an argument of the utility function, and, at the same time, the dominant asset in the portfolio. This paper uses a mean-variance efficiency framework to examine the household's optimal portfolio problem when owner-occupied housing is included in the list of available assets. Housing differs from stocks and bonds in a crucial way: since the household's ownership of residential real estate determines the level of its consumption of housing services, the household's demand for real estate is "overdetermined" in the sense that the level of real estate ownership which is optimal from the point of view of the consumption of housing services may differ from the optimal level of housing assets from a portfolio point of view. With rental markets for housing, a household can, in principle, divorce the size of its holdings of real estate assets from the level of housing services it consumes. However, rental housing is by no means a perfect substitute for owner-occupied housing. We assume, instead, that the preferential tax treatment of owner-occupied housing and the transactions costs and agency costs involved in the rental market for housing create frictions large enough to effectively constrain the household to include in its asset portfolio the level of housing consistent with its consumption of housing services.¹ The paper focuses on the impact of the portfolio constraint imposed by the

consumption demand for housing on the household's optimal holdings of financial assets.

Section II of the paper is similar in spirit to a recent paper by Jan K. Brueckner (1997), which analyzes the interaction between the consumption demand and the investment demand for housing in a mean-variance portfolio model. Brueckner considers a general covariance matrix and mean vector of returns and a general utility function, and derives analytical results. In contrast, our implementation of the mean-variance framework is quantitative. That is, we estimate the covariance matrix and vector of expected returns for housing and financial assets and solve for the efficient frontiers and optimal portfolios numerically.

The risk characteristics of housing are estimated using two distinct sources: data from the Panel Study of Income Dynamics (PSID) and data from Karl E. Case and Robert J. Shiller (1989) based on repeat sales transactions prices for four U.S. cities. Both data sources indicate that housing prices have a large idiosyncratic component; the standard deviation of the return to housing, at the level of the individual house, is about 0.14. In addition to housing, the portfolio can include nonnegative amounts of Treasury bills, Treasury bonds, and stocks. The household can borrow only in the form of a mortgage, which is limited to 100 percent of the value of the house. Using the estimated vector of expected returns and covariance matrix of asset returns, we plot the constrained mean-variance efficient frontiers for various values of the household's ratio of house value to wealth, h (the "housing constraint"). The housing constraint has an enormous effect on the risk and return trade-off available to the household. Young households, which typically have large holdings of real estate relative to their net worth, are highly leveraged and therefore forced into a situation of high risk (and return). As a result, these young households have a strong incentive to reduce the risk of their portfolio by using their net worth to either pay down their

* Flavin: Department of Economics, University of California, San Diego, CA 92093, and National Bureau of Economic Research (e-mail: mflavin@ucsd.edu); Yamashita: Department of Economics, University of Nevada, Las Vegas, NV 89154 (e-mail: tyamashi@ccmail.nevada.edu). We thank Elena Bisagni, Jan Brueckner, Wouter den Haan, James Hamilton, Bruce Lehmann, Greg Mankiw, and the referees for comments, and Robert Shiller for providing the house price transactions data.

¹ The implications of the dual role of housing (as both a consumption good and an investment good) for tenure decisions of households were first analyzed by J. Vernon Henderson and Yannis M. Ioannides (1983).

mortgage or buy bonds instead of buying stocks. In contrast, the optimal portfolio share of stocks is larger for the lower ratios of housing to net worth typical of older households.

In addition to the implications for the composition of the portfolio over the life cycle, the inclusion of housing has implications for cross-sectional scatters of portfolio data. In the final section, we use the model to offer a potential resolution to the “asset allocation puzzle” posed by Niko Canner et al. (1997). Using as data points the portfolio shares recommended by financial advisors in the popular press for investors of varying degrees of risk aversion, Canner et al. point out that financial advisors tend to recommend that the ratio of bonds to stocks increase with the level of risk aversion, in apparent contradiction of the mutual fund separation theorem. Compared to a world in which the list of assets is restricted to the standard financial assets (T-bills, bonds, and stocks), the inclusion of housing has the effect of altering the risk and return trade-off in such a way that most households are at a corner with respect to T-bills. That is, the nonnegativity constraint on T-bills is binding for a wide range of values of h and of the degree of risk aversion. Since households are at a corner with respect to the riskless asset, different degrees of risk aversion cannot be accommodated by varying the relative quantities of the riskless asset and the risky portfolio; instead, preferences toward risk are accommodated by varying the ratio of bonds to stocks.²

I. Housing in a Mean-Variance Efficiency Framework

Of the extensive literature on efficient portfolios, only a few papers incorporate real estate as an asset. William N. Goetzmann and Roger Ibbotson (1990) and Goetzmann (1993) used regression estimates of real estate price appreciation, and Stephen A. Ross and Randall C. Zisler (1991) calculated returns from real estate investment trust funds, to characterize the risk

and return to real estate investment. While regionally diversified real estate funds have recently become available, the vast majority of households invest in real estate by purchasing a particular house, rather than by purchasing shares of a diversified real estate fund. Further, the returns to an investment in a real estate fund and the returns to an investment in one’s personal residence receive very different tax treatment. In analyzing the efficient frontier facing a large institutional investor such as TIAA-CREF (Teachers Insurance and Annuity Association-College Retirement Equities Fund), Ross and Zisler have appropriately used diversified real estate funds to characterize the risk characteristics of real estate.

Our focus, in contrast, is on the portfolio problem faced by the typical household. We assume that households hold real estate in the form of a specific house (rather than investing in a diversified fund and renting to satisfy their demand for housing services) due to tax distortions and transactions or agency costs associated with renting, but do not explicitly model the renting versus owning decision. We abstract from labor income or human wealth, and assume that wealth is held only in the form of financial assets and housing. The household can invest in any of n risky financial assets, but there is no riskless asset. Further, the household can borrow up to the value of the house in the form of a mortgage; other than the mortgage, all financial assets must be held in nonnegative amounts. As in Sanford J. Grossman and Guy Laroque (1990), we assume that once the household purchases a particular house, no adjustments to the size (or any other attribute such as location) can be made without selling the existing house, incurring an adjustment cost proportional to the value of the house, and buying a new house. Unlike the house, holdings of financial assets (including the mortgage) can be adjusted with zero transactions cost.

The household’s total wealth, W_t , is given by:

$$(1) \quad W_t = \mathbf{X}_t \ell + P_t H_t$$

where $\mathbf{X}_t = (1 \times n)$ vector of amounts (expressed in terms of a nondurable consumption good used as the numeraire) held of the n risky assets, $\ell = (n \times 1)$ vector of ones, $H_t =$ the physical quantity held of housing,

² In this paper, we show the importance of owner-occupied housing in explaining the cross-sectional and temporal variations in portfolio composition. For effects of other sources of background risks, such as entrepreneurial risks and labor-income risks, see John Heaton and Deborah Lucas (2000).

measured in square feet, and P_t = the price per square foot of housing relative to the price of the nondurable good. Using the last element of \mathbf{X}_t to represent the mortgage, the corner constraints on the vector of financial assets are given by:

$$(2) \quad -P_t H_t \leq X_{n,t} \leq 0$$

(constraint on mortgage borrowing)

$$(3) \quad 0 \leq X_{i,t} \quad i = 1 \text{ to } n - 1$$

(nonnegativity constraints on other financial assets)

where $X_{i,t}$ is the i th element of \mathbf{X}_t . Note that because the household can borrow only in the form of a mortgage, it must hold a mortgage if the value of the house exceeds net worth.

Denote the real, after-tax return on financial asset i in year t as $R_{i,t}$ and the real, after-tax return on housing as $R_{H,t}$; the asset return is the sum of the expected return (μ_i or μ_H) and a stochastic component:

$$(4) \quad R_{i,t} = \mu_i + \varepsilon_{i,t}$$

$$R_{H,t} = \mu_H + \varepsilon_{H,t}.$$

The covariance matrix of returns is denoted by the $(n + 1)$ by $(n + 1)$ matrix $\mathbf{\Omega}$:

$$(5) \quad \mathbf{\Omega} \equiv E(\boldsymbol{\varepsilon}_t \boldsymbol{\varepsilon}_t^T)$$

where $\boldsymbol{\varepsilon}_t$ is the $((n + 1) \times 1)$ vector $\boldsymbol{\varepsilon}_t \equiv [\varepsilon_{1,t}, \dots, \varepsilon_{n,t}, \varepsilon_{H,t}]^T$. Both the vector of expected returns, $\boldsymbol{\mu}$, ($\boldsymbol{\mu} \equiv [\mu_1, \mu_2, \dots, \mu_n]^T$) and the covariance matrix of returns, $\mathbf{\Omega}$, are assumed time invariant.

Dividing equation (1) by W_t , the household's asset holdings can be restated as shares of total wealth rather than as levels:

$$(6) \quad 1 = h_t + \mathbf{x}_t \boldsymbol{\ell}$$

where $h_t \equiv \frac{P_t H_t}{W_t}$ and $\mathbf{x}_t \equiv \frac{\mathbf{X}_t}{W_t}$.

The optimal portfolios calculated in Section II are based on the assumption that the household maximizes a function of the mean and

variance of the return to its asset portfolio, inclusive of housing.³ Stating both the maximization problem and the constraints in terms of shares of wealth, the household's problem is:

$$(7) \quad \max_{\mathbf{x}} \left\{ (\mathbf{x}_t \boldsymbol{\mu} + h_t \mu_H) - \frac{\Lambda}{2} [\mathbf{x}_t, h_t] \mathbf{\Omega} [\mathbf{x}_t, h_t]^T \right\}$$

subject to:

$$1 = h_t + \mathbf{x}_t \boldsymbol{\ell}$$

(adding up constraint)

$$-h_t \leq x_{n,t} \leq 0$$

(constraint on mortgage borrowing)

$$0 \leq x_{i,t} \quad i = 1 \text{ to } n - 1$$

(nonnegativity constraints on other financial assets).

In essence, we model the household as maximizing over its holdings of financial assets, conditional on the current value of the state variable which represents the ratio of house value to net worth, h_t . While we do not formally model the house purchase decision, we follow Grossman and Laroque (1990) in assuming that at every moment, the household considers whether the disparity between the current-size house and the frictionlessly optimal-size house is sufficiently large to justify paying the transactions cost and reoptimizing over the house, H_t . If it is optimal to sell the house immediately, the household (instantaneously) pays the transactions cost, sells the old house, and buys a new house. Thus the home purchase decision is endogenous and fully rational, but, because of the transactions cost, infrequent. For example, in the model studied by Grossman and Laroque (1990), a conservative estimate of the transactions cost of 5 percent of the value of the house sold implies that the average time between house purchases is 20 to 30 years.

Since the state variable h_t places a constraint

³ The issue of whether, and under what conditions, the mean-variance optimization framework can be obtained as the outcome of the household's lifetime utility maximization problem is not addressed in this paper.

on the optimization problem, we refer to the value of h_t as “the housing constraint.” For a given value of h_t , the mean-variance efficient frontier available to the household can be calculated by finding the value of \mathbf{x}_t which achieves the minimum variance portfolio for a given expected return. Thus the constrained mean-variance efficient frontier available to the household depends of the value of the state variable h_t . In equation (7), the household’s preferences toward risk are reflected in the parameter A ; higher values of A indicate a greater degree of risk aversion. For the preferences expressed in (7), the slope of the household’s indifference curve is:

$$(8) \quad \frac{\partial \bar{\mu}}{\partial \sigma} = A\sigma$$

$$\text{where} \quad \bar{\mu} = \mathbf{x}_t\boldsymbol{\mu} + h_t\mu_H$$

$$\sigma^2 = [\mathbf{x}_t, h_t]\boldsymbol{\Omega}[\mathbf{x}_t, h_t]^T.$$

Because of the corner constraints, the optimization problem is solved by numerical methods; the optimal portfolio will depend both on the state variable h_t and on preferences toward risk, A . To estimate returns on housing, we use data from the 1968–1992 waves of the PSID. Every year, the PSID asks homeowners how much their house would sell for if the house were put on the market on the date of the interview, enabling us to calculate the return to owner-occupied housing at the household level.⁴ The return to housing depends on appreciation of the value of the house, the value of the flow of housing services, and costs of ownership and maintenance. Lacking direct observations on the rental value of the house and maintenance costs, these components are modeled as follows:

$$(9) \quad D_t = (r + d)P_{t-1} + \text{PropertyTax}_t$$

$$(10) \quad \text{COM}_t = dP_{t-1} + (1 - \tau)\text{PropertyTax}_t$$

⁴ As a rough measure of the accuracy of subjective housing value measures, Jonathan Skinner (1994) compares the annual rate of self-reported price changes with the objective Commerce Department measures and finds the two series are quite close in mapping house-price changes over the 1970’s and 1980’s.

where r is the real interest rate, d is the depreciation rate, and τ is the marginal income tax rate. The imputed annual rental value (analogous to the dividend on a stock), denoted D_t , reflects the assumption that property taxes are passed through into rents.⁵ In the absence of expenditures on maintenance and repairs, physical depreciation (at rate d) would be reflected in the real value of the house (P_t). However, we assume that both landlords and homeowners spend on maintenance and repairs an amount equal to the annual depreciation of the house so that the physical condition of the house is constant. In addition to maintenance and repairs, the cost of ownership and maintenance (COM_t) includes the net property tax payment (i.e., net of the deduction against income taxes). The real return on housing, $R_{H,t}$, is:⁶

$$(11) \quad R_{H,t} = \frac{P_t + D_t - \text{COM}_t - P_{t-1}}{P_{t-1}} \\ = \frac{P_t + rP_{t-1} + \tau\text{PropertyTax}_t - P_{t-1}}{P_{t-1}}.$$

⁵ In equation (9), we are modeling the market rent which the homeowner would pay if he explicitly rented the house from a landlord rather than owning it himself. Equation (9) does not contradict the public finance result that states that the incidence of the property tax falls on the landlord in the form of lower property values, and cannot be shifted onto renters. Since the house price data reflects the market value of the house under the existing tax regime, as opposed to the value the house would have in the absence of a property tax, the fact that the landlord “pays” the tax in the form of lower property values is already incorporated in the house price data. Given data on the market value of a particular house, equation (9) is a simple model, based on the zero-profit condition, of how the landlord would set the market rent on that particular house; i.e., the market rent would be sufficient to cover the landlord’s costs (interest or forgone interest, depreciation, and the property tax bill).

⁶ A possible bias arises from the fact that the house price is not the actual transaction price but the owner’s estimate. If there is a systematic bias in the owner’s estimate of the house price, then the return calculation may also be biased. Using the panel aspect of the 1985 and 1987 waves of the American Housing Survey, John L. Goodman, Jr. and John B. Ittner (1993) compare owners’ estimates with subsequent sales prices and find that the mean error of the homeowner’s estimate is 8.3 percent, and that the errors are not correlated with characteristics of the house or demographic characteristics of the owners. Given Goodman and Ittner’s estimate of the homeowners’ bias, it is straightforward to show, using equation (11), that the mean overestimation of 8.3 percent of the real house price would give us a downward bias in the calculated rate of return of around 0.06 percent, assuming a 33-percent marginal tax rate and a 2.5-percent property tax rate.

In computing the real return to housing, the nominal house value and nominal property tax payments as reported by the respondent are converted into real terms using the CPI-U deflator to obtain P_t and $PropertyTax_t$. The short-term interest rate, r_t , is assumed to be a fixed 5 percent. In general, the marginal income tax rate, τ , will vary both cross-sectionally and across time. For most waves of the PSID, it is possible to impute the household's marginal tax rate; thus, one could construct a series of real returns to housing using a household- and time-specific value of the marginal tax rate. However, it is not necessary to interpret "the household" as a representative household (i.e., one of many identical households) or as an average household, since the analysis is entirely partial equilibrium and no aggregate market-clearing conditions are imposed. Instead, "the household" is interpreted as a specific, hypothetical household which faces a known marginal income tax rate of 33 percent (28 percent federal and 5 percent state).⁷

In order to incorporate tax effects, all asset returns are stated in after-tax, real terms. In the case of the return to housing, the imputed rent portion of the return is not taxed. While capital gains on housing are taxable in some circumstances, we assume that the bulk of realized capital gains escape taxation due to rollover provisions and the fact that inherited real estate receives a stepped-up cost basis when transferred. We therefore assume that neither the rental value nor the capital gains portion of the

return is taxed, and interpret $R_{H,t}$ as the after-tax real return to housing.

In calculating real after-tax mortgage rates, we assume that the household takes out a fixed-rate mortgage in the year in which the house is purchased and does not refinance as long as they remain in the same house. Taking into account the fact that the nominal rather than real mortgage interest payments are deductible, the household pays a real after-tax interest rate of:

$$(12) \text{ Mortgage}_t = \frac{1 + (1 - \tau) \text{Nominal Mortgage}_s}{1 + \text{Inflation}_t} - 1.$$

In equation (12), $\text{Nominal Mortgage}_s$ denotes the fixed rate at which the mortgage was issued when the home was purchased in year s and Mortgage_t denotes the after-tax real interest rate on the mortgage in a subsequent year t . In implementing equation (12), $\text{Nominal Mortgage}_s$ is measured by the annual average of conventional home mortgage rates charged by major lenders in year s , τ is the hypothetical marginal tax rate of 33 percent, and the inflation rate is measured by the CPI-U. The household is not bound by a fixed amortization schedule, but instead has flexibility in optimizing over the size of the mortgage. For example, the household can reduce its degree of leverage by paying down the mortgage balance. Conversely, the household can increase its total mortgage indebtedness by taking out a second mortgage without having to refinance the original mortgage.

The return on Treasury bills is measured by constructing the holding period return generated by rolling over short-term Treasury bills for a period of one year, assuming that the holding period return is taxed at a marginal rate (federal only) of 28 percent, then converting the after-tax nominal rate into real terms. Similarly, the return to Treasury bonds is measured by deflating the after-tax (28-percent) holding period yield on 20-year Treasury bonds.

Data on the S&P 500 is used to measure stock returns. Using separate data series on dividend income and capital gains, we assume that dividend income is taxed at the rate of 33 percent, while the capital gains portion of the return escapes taxation completely. The after-tax nominal

⁷ An alternative approach would be to estimate the expected return vector and covariance matrix of after-tax returns by taking sample averages of the historical data on after-tax returns. In this approach, the data on after-tax returns would reflect variation over time in marginal tax rates as well as variation in pretax returns. Because it implicitly assumes that the household faces a marginal tax rate which is a random draw from the historical distribution of marginal tax rates, we think that this approach overstates the degree of randomness in the marginal tax rate facing the household at a particular point in time. However, as a robustness check, we calculated an alternative version of Table 1A in which historical variation over time of the federal marginal tax rate is included. That is, for each year, we calculated the marginal federal tax rate of the PSID household that ranked in the 75th percentile for income, and used the year-specific marginal federal tax rate to calculate after-tax real returns. A comparison of these estimates, reported in the Data Appendix, with the corresponding estimates in Table 1A, indicates that the alternative method of incorporating taxes makes very little difference.

TABLE 1A—EXPECTED RETURNS AND COVARIANCE MATRIX—PSID DATA

	T-Bills	Bonds	Stocks	Mortgage	House
Mean return (arithmetic)	-0.0038	0.0060	0.0824	0.0000	0.0659
Standard deviation	0.0435	0.0840	0.2415	0.0336	0.1424
Covariance Matrix					
T-bills	0.0018920				
Bonds	0.0025050	0.0070613			
Stocks	0.0002008	0.0040381	0.0583292		
Mortgage	0.0007087	0.0023854	0.0025400	0.0011274	
House	-0.000119	-0.000067	-0.000178	-0.0000057	0.020284
Correlation Matrix					
T-bills	1.0000				
Bonds	0.68533 (0.09103)	1.0000			
Stocks	0.01912 (0.12498)	0.19897 (0.12251)	1.0000		
Mortgage	0.84119 (0.11529)	0.680286 (0.15626)	0.467954 (0.18842)	1.0000	
House	-0.03339 (0.21309)	-0.004506 (0.21320)	-0.000771 (0.21319)	-0.001192 (0.21320)	1.0000

Note: Standard errors are in parentheses.

return is then converted to a real return. Note that by assuming that capital gains are not taxed, our measure of the after-tax return to stocks probably overstates the actual return.

Real after-tax mortgage rates and returns to housing are calculated for homeowners in the PSID from 1968 to 1992. Excluding the poverty subsample, all households that owned a house over at least a two-year period between 1968 and 1992 were included (1,817 households). For each household, the return to housing $R_{H,t}$ was calculated using equation (11) if the household was living in the same house during year t and year $t - 1$. If, in year t , the household moved to a new house or became a renter, the value of $R_{H,t}$ was set to missing.

The calculation of returns to housing relies on the change in the reported value of the house between interviews, which are usually conducted between February and April. We interpret the housing return series as measuring the return over the March-to-March interval, and construct the returns on financial assets for the same time intervals. While the data on housing returns and mortgage rates are limited to the 1968–1992 period, data for a longer period (1926–1992) are used for T-bills, T-bonds, and the S&P 500. A detailed description of the data

and calculation of rates of return is provided in the Data Appendix.

The mean returns, standard deviations, and covariance matrix of the five assets are reported in Table 1A. In calculating the sample statistics of mean returns on housing assets and their covariance with other financial assets, we assume that all returns are drawn from the same distribution regardless of region and year. As a robustness measure, we trimmed the top and bottom 2 percent of housing returns on the grounds that the tails contain some extreme values likely to reflect gross measurement errors. As expected, stocks have the highest after-tax real rate of return at 8.2 percent, and Treasury bills the lowest, with negative 0.4 percent. The real return to housing at 6.6 percent is comparable to that of stocks, while the return to housing has a substantially smaller standard deviation than the stock return. The after-tax real interest rate on a fixed mortgage is on average zero over this period. According to the bottom row of the correlation matrix, the return to housing is uncorrelated with the return to each of the financial assets. 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 in

TABLE 1B—EXPECTED RETURNS AND COVARIANCE MATRIX—CASE-SHILLER PRICE INDICES

	T-Bills	Bonds	Stocks	Atlanta	Chicago	Dallas	San Francisco
Mean return	-0.0038	0.0060	0.0824	0.05356	0.05363	0.07196	0.09787
Standard deviation	0.0435	0.0840	0.2415	0.04200	0.06079	0.04872	0.06540
Covariance Matrix							
T-bills	0.0018920						
Bonds	0.0025050	0.0070613					
Stocks	0.0002008	0.0040381	0.0583292				
Atlanta	0.0005253	0.0020381	0.003202	0.001764			
Chicago	0.0002768	0.002859	0.002211	0.001006	0.003696		
Dallas	-0.000127	-0.000769	-0.000825	0.000851	0.001327	0.002373	
San Francisco	-0.00058	0.000415	-0.000223	-0.000159	0.001931	0.000796	0.004277
Correlation Matrix							
T-bills	1.0000						
Bonds	0.68533 (0.09103)	1.0000					
Stocks	0.01912 (0.12498)	0.19897 (0.12251)	1.0000				
Atlanta	0.41871 (0.24271)	0.38527 (0.24663)	0.42041 (0.23970)	1.0000			
Chicago	0.15244 (0.26414)	0.37332 (0.24794)	0.20051 (0.26001)	0.39412 (0.24563)	1.0000		
Dallas	-0.08701 (0.26625)	-0.12527 (0.26516)	-0.09341 (0.26632)	0.41585 (0.24306)	0.44796 (0.23895)	1.0000	
San Francisco	-0.29702 (0.25520)	0.05041 (0.26692)	-0.01879 (0.26721)	-0.05794 (0.26681)	0.48567 (0.23362)	0.24987 (0.25878)	1.0000

Note: Standard errors are in parentheses.

terms of numerical size and statistical significance.⁸

In many ways the PSID data is ideally suited to our purpose; it is a large, nationally representative sample and measures house prices at the level of the individual house, with no geographic aggregation. Unlike transactions price data, which may be distorted by composition effects due to variation over time in the transactions volume of different segments of the real estate market, the PSID tracks the same households over time. However, the fact that the PSID house price data is based on the owner's valuation rather than on transactions prices is a potential problem. As a check on the PSID data, the covariance matrix was also estimated using data from Case and Shiller (1989) based on repeat sales transactions prices for four cities—

Atlanta, Chicago, Dallas, and San Francisco. Table 1B reports expected returns and covariances analogous to Table 1A, except that the housing return is calculated from an index of housing prices constructed by Case and Shiller for each of the four cities separately.⁹ The expected return to housing in the four cities is 0.054 (Atlanta), 0.054 (Chicago), 0.072 (Dallas), and 0.098 (San Francisco), compared to 0.066 for the national sample in the PSID. For none of the four cities is the correlation between the housing return and financial asset returns

⁹ Using data from the Society of Real Estate Appraisers, Case and Shiller (1989) construct a Weighted Repeated Sales Index for each of the cities by extracting the date (year and quarter) and sales price of houses which sold twice during the sample period (1970–1986). Using weighted least squares, the city index is estimated by regressing the change in the log of the individual house price on a set of dummy variables. With this method, they can estimate an index for the average house appreciation in each city, as well as assess the magnitude of idiosyncratic house price changes around the citywide index.

⁸ The finding that housing returns are essentially uncorrelated with stock returns is consistent with the findings of Goetzmann (1993) using repeated sales regression estimates.

statistically significant. The lower right block of the correlation matrix reports the point estimates and standard errors of the correlation between returns to housing in different cities. With the partial exception of San Francisco, the correlation between the housing return in two different cities is positive and approaching statistical significance at the 5-percent level. For all city pairs, the correlation of housing returns is statistically significantly different from unity, and for most pairs is about 0.4.

Because the housing return series in Table 1B are based on citywide price indices, instead of returns to individual house prices, the estimated standard deviation of housing returns is substantially smaller (ranging from 0.042 to 0.065) than the standard deviation of 0.142 for the PSID data. In their analysis of individual house prices, Case and Shiller (1989) report that in a regression of annual real individual house price changes on the contemporaneous annual change in the citywide index, “the R^2 is only 0.066 for Atlanta, 0.158 for Chicago, 0.121 for Dallas, and 0.270 for San Francisco,”¹⁰ indicating that individual house prices have an important idiosyncratic component apart from the city-specific component. Case and Shiller conclude that “individual house prices are like many individual corporate stock prices in the large standard deviation of annual percentage change, close to 15 percent a year for individual housing prices.”¹¹

Data from both sources are consistent with the following conclusions: the return to housing is uncorrelated with the returns to financial assets, and the standard deviation of housing returns at the level of the individual house is about 0.14 to 0.15. Since the PSID data reflects a nationally representative sample and covers a longer time period than the Case-Shiller data, we use the covariance matrix and vector of expected returns estimated from the PSID to calculate the efficient frontier and portfolios.¹²

¹⁰ Case and Shiller (1989 p. 127).

¹¹ Case and Shiller (1989 p. 127).

¹² Because the estimates of the covariance of the return to housing with the returns to financial assets are so close to zero, in terms of magnitude as well as statistical significance, we impose the assumption that the housing return is uncorrelated with the return to financial assets by substituting values of zero in place of the point estimates in the first four elements of the last row of the covariance matrix reported in Table 1A.

TABLE 2—LIFE-CYCLE PATTERN OF ASSET HOLDINGS—1989 PSID WEALTH DATA (MEAN RATIO OF ASSET TO NET WORTH)

Age of Head	Cash	Bonds	Stocks	House	Mortgage
18–30	0.193	0.072	0.056	3.511	–2.833
31–40	0.169	0.067	0.068	2.366	–1.671
41–50	0.148	0.060	0.085	1.588	–0.882
51–60	0.200	0.058	0.092	0.969	–0.319
61–70	0.254	0.048	0.113	0.757	–0.171
71+	0.264	0.029	0.098	0.648	–0.038

Source: Authors’ tabulation from the PSID.

II. Optimal Portfolios and the Constrained Efficient Frontier

In our approach, the ratio of house value to wealth, h , is treated as an endogenously determined state variable which imposes a constraint on the mean-variance optimization problem. The value of h varies dramatically over the life cycle, as documented in Table 2, which reports the mean asset-to-net worth ratio, by age-group, of PSID homeowners in 1989. On average, homeowners in the youngest cohort (aged 18–30) are highly leveraged, with a house-to-net wealth ratio of 3.51. As the household accumulates wealth, the value of h drops to approximately unity for the cohort aged 51–60 and to $h = 0.65$ for the oldest homeowners. Note that the data on the cohort mean asset shares is consistent with the model’s constraint on mortgage borrowing ($-h \leq x_{n,t} \leq 0$) and that the life-cycle variation in h is reflected in an equally dramatic pattern in the ratio of mortgage to net wealth.

In Table 2, “Cash” refers to any money held in checking and savings accounts, money market accounts, or Treasury bills, and “Bonds” refers to any other savings including bond funds and other assets such as rights in a trust and cash values of insurance policies. Thus these two categories correspond only very loosely to the Treasury bills or 20-year Treasury bonds used in the estimation of the risk characteristics of financial assets.

In Figure 1 we plot the constrained mean-variance efficient frontier of the portfolio, inclusive of housing, which corresponds to each of six different values of the ratio of house value to net worth ranging from 3.51 (the average value of h among PSID homeowners 18–30 years




FIGURE 1. EFFICIENT FRONTIERS WITH DIFFERENT HOUSING CONSTRAINTS

old) to 0.65 (the value of h typical of homeowners 71 and older). That is, for each of the stated values of h , quadratic programming was used to solve for the vector of financial assets, \mathbf{x} , which satisfies the constraints in (7) and achieves the minimum-variance portfolio for a given expected return. The constrained efficient frontier for $h = 0$ is also plotted.

From among the efficient portfolios, the optimal portfolio is identified by the point of tangency of the constrained frontier with the slope of the indifference curve, equation (8). Since the slope of the indifference curve is a function of the household's degree of risk aversion, A , we report, in Table 3, the optimal portfolios corresponding to various combinations of the constraint, h , and of risk aversion, as measured by A . A subset of the optimal portfolios (i.e., those for values of $A = 2, 4$, or 8) are plotted in Figure 1. Also plotted are the portfolios corresponding to the average asset shares for each PSID age cohort in Table 2 (represented by x 's, in Figure 1).

For very low levels of risk aversion ($A \leq 1$), the optimal portfolio is at the upper corner of the constrained frontier for the whole range of values of h . These risk-tolerant households maximize their expected return by holding a 100-percent mortgage on their house, and allocating 100 percent of their financial asset portfolio to stocks. For $A = 2$, the optimal portfolio

of financial assets consists of about 72 percent stocks, 28 percent bonds, and no T-bills. For both $A = 2$ and $A = 1$, the house is fully mortgaged and the optimal shares of financial assets do not vary noticeably with the state variable, h . For moderate and high levels of risk aversion (that is, for values of A ranging from 4 to 10), the household's portfolio of financial assets shows a clear life-cycle pattern. For example, for a moderate level of risk aversion ($A = 4$), the household remains fully mortgaged for high levels of h , but starts paying off the mortgage when h falls below 1.5. For households with the high levels of h which are typical of young households, the optimal portfolio of financial assets consists of 60 percent bonds, 40 percent stocks, and no T-bills. When the value of h falls to one or lower, as is characteristic of households aged 50 or over, the optimal portfolio includes a larger share of stocks and a smaller share of bonds. For example, the optimal portfolio of a household with a housing to net worth ratio of $h = 0.65$ would consist of 44 percent bonds and 56 percent stocks. An amplified version of the same life-cycle pattern holds for the higher levels of risk aversion ($A = 8$ or 10). These more highly risk-averse households start paying down their mortgages almost immediately; by the time the housing to net worth ratio falls to 0.65, the mortgage is fully paid off. For high levels of h (i.e., $h = 3.51$), the

TABLE 3—OPTIMAL PORTFOLIO WEIGHTS FOR DIFFERENT CONSTRAINTS ON h

Housing-to- NW Ratio	Assets in Portfolio	Degree of Risk Aversion, A				
		A = 1	A = 2	A = 4	A = 8	A = 10
$h = 0$	Treasury bills	0	0	0.4723	0.7931	0.8279
	Treasury bonds	0	0.2800	0.1527	0	0
	Stocks	1	0.7200	0.3750	0.2069	0.1721
$h = 3.51$	Treasury bills	0	0	0	0	0
	Treasury bonds	0	0.2728	0.6001	0.7703	0.8096
	Stocks	1	0.7272	0.3999	0.2297	0.1904
$h = 2.37$	Mortgage	-1	-1	-1	-1	-1
	Treasury bills	0	0	0	0	0
	Treasury bonds	0	0.2705	0.6109	0.7603	0.7870
$h = 1.59$	Stocks	1	0.7295	0.3891	0.2397	0.2130
	Mortgage	-1	-1	-1	-0.9724	-0.9327
	Treasury bills	0	0	0	0	0
$h = 0.97$	Treasury bonds	0	0.2780	0.6053	0.6763	0.6958
	Stocks	1	0.7220	0.3947	0.3238	0.3042
	Mortgage	-1	-1	-1	-0.7630	-0.7186
$h = 0.76$	Treasury bills	0	0	0	0	0
	Treasury bonds	0	0.2795	0.5164	0.5103	0.5064
	Stocks	1	0.7205	0.4836	0.4897	0.4936
$h = 0.65$	Mortgage	-1	-1	-0.7622	-0.3737	-0.2766
	Treasury bills	0	0	0	0	0
	Treasury bonds	0	0.2815	0.4674	0.3947	0.3757
$h = 0.65$	Stocks	1	0.7185	0.5326	0.6053	0.6243
	Mortgage	-1	-1	-0.6111	-0.0844	0
	Treasury bills	0	0	0	0.1357	0.3744
$h = 0.65$	Treasury bonds	0	0.2782	0.4360	0.3217	0.2018
	Stocks	1	0.7218	0.5640	0.5426	0.4237
	Mortgage	-1	-1	-0.4689	0	0

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

strongly risk-averse households have portfolios of financial assets consisting of about 80 percent bonds and 20 percent stocks. For low levels of h (i.e., $h = 0.65$), the portfolio consists of about half stocks, and half nominal assets.¹³

The constraint that the value of the mortgage cannot exceed the value of the house is often

binding in the optimal portfolios depicted in Figure 1. For example, consider the constrained efficient frontier for $h = 3.51$. Along the entire length of the constrained frontier, the household holds a mortgage equal in value to the house and holds an amount equal to its net wealth in a combination of stocks and bonds. As risk aversion declines, and the household moves along the constrained frontier toward riskier portfolios, this movement is achieved by reducing the bond share and increasing the stock share while the mortgage remains at the maximum value consistent with the constraint in equation (7). Thus the slope of the constrained frontier reflects the risk and return trade-off involved in trading bonds for stocks. Once the portfolio consists of $h = 3.51$ with a 100-percent mortgage, stocks = 1, bonds = 0, and T-bills = 0, the fact that the mortgage constraint is binding implies that the household has obtained the highest return portfolio which is consistent with $h = 3.51$.

¹³ Michael Fratantoni (1996) studies the effects of homeownership on risky asset holding by modeling the commitment to a fixed mortgage payment combined with uncertain labor income in a simulation model. In his model, the share of risky assets in total wealth is higher for young households who have just bought a house than for older households nearing retirement. João F. Cocco (1998) also uses simulation to solve a life-cycle model with risky labor income for the optimal holding of housing and financial assets. He finds that compared to a model without housing, the introduction of housing increases the proportion of financial assets held in the form of stocks later in life. Nevertheless, stock holding as a proportion of financial assets is predicted to decline over the life cycle.

Given the estimated vector of expected returns and covariance matrix, the model predicts that households generally hold both a mortgage and bonds. Since mortgages and bonds are similar assets, an alternative approach would be to consider a mortgage and a bond as perfect substitutes; i.e., assume that there is a single financial instrument which is called a mortgage if one takes a negative position and called a bond if one takes a positive position, in addition to stocks, T-bills, and housing. Under this alternative approach to modeling the mortgage, the model can give predictions about the sum of holdings of mortgage and bonds, but cannot say anything about the shares of mortgage and bonds separately. When the portfolios are calculated for this simplified menu of assets (i.e., a mortgage is a negative position in bonds), the basic characteristics of the portfolios in Table 3 are confirmed: the optimal holding of T-bills is at the corner unless h is low and risk aversion is high, the degree of leverage is reduced as h falls over the life cycle, more risk-averse households reduce their degree of leverage more rapidly than less risk-averse households, and the percentage of net wealth held in the form of stocks increases as h falls.

III. Life-Cycle and Cross-Sectional Implications

At its most basic level, the model implies that the efficient frontier faced by a household is not static, but varies with the state variable h . Since h varies over the life of a household, the composition of the optimal portfolio of financial assets in general will vary over the life cycle, even if preferences and the stochastic structure of the asset markets are constant. The basic message for the life cycle of portfolio composition is that young households with high values of h are forced by their high degree of leverage to hold a risky portfolio and therefore use their net worth to reduce portfolio risk rather than attempting to further increase their expected return. To reduce portfolio risk, the young households hold only a small to moderate share (depending on the degree of risk aversion) of stocks, and reduce their leverage, either explicitly by paying down the mortgage, or implicitly by keeping the 100-percent mortgage and simultaneously holding bonds. As wealth accumulates and the value of h falls, the household shifts a greater fraction of its financial assets

into stocks. The percentage of financial assets devoted to stocks peaks at about 60 percent when h is in the range $0.4 < h < 0.8$, then declines for lower levels of h .

Because the percentage of financial assets devoted to stocks is inversely related to h over most of its range, and h tends to decline as the household ages, the model suggests that stock holding will increase with age. Existing studies that have tested for changes in the composition of household portfolios over the life cycle have found conflicting evidence of age effects. Using data from the 1978 Survey of Consumer Financial Decisions, Mervyn A. King and Jonathan I. Leape (1987) find that the probability of owning stocks, corporate bonds, municipal bonds, and savings bonds increases with age, while the probability of owning Treasury bonds declines with age. Ioannides (1992), however, using the 1983 and 1986 Surveys of Consumer Finances, finds that coefficients on age are not significant at the 5-percent level in probit regressions of ownership of most financial assets, except for IRAs, housing equity, and other debt. Carol C. Bertaut (1996), using the 1983–1989 panel of the Survey of Consumer Finances, finds that the probability of owning stocks increases with age. Using the 1983, 1989, and 1992 Surveys of Consumer Finances, James M. Poterba and Andrew A. Samwick (1997) find strong age effects of both the probability of ownership and the portfolio shares of stocks (directly held as well as all taxable equity, including brokerage accounts and equity mutual funds). Yamashita (2001), using the 1989 Survey of Consumer Finances, however, finds that the coefficient on age becomes insignificant in estimating the share of stocks in the portfolio, when the housing constraint is included in the specification. Using various waves of the Survey of Consumer Finances along with a panel data set of TIAA-CREF participants over 1987–1999, John Ameriks and Stephen P. Zeldes (2000) find that portfolio shares of stocks increase strongly with age. Examining those households that owned at least some equities in their portfolio, however, Ameriks and Zeldes find a nearly flat age profile in both the SCF and the TIAA-CREF data, implying an absence of age effects.

The inclusion of housing also has implications for the composition of portfolios in a cross section. According to Table 3, a model which

ignores housing would predict that households with moderate to high risk aversion ($A \geq 4$) will hold substantial quantities of T-bills (accounting for as much as 47 percent of financial assets for $A = 4$ and 83 percent of financial assets for $A = 10$). If housing is included in the portfolio, however, the constraint that the household cannot borrow at the T-bill rate becomes binding for a large range of values of h and A . That is, for the entire region $h \geq 0.76$ and $A \leq 10$, the household is at a corner with zero T-bills. Only when h falls to 0.65 in conjunction with high levels of risk aversion ($A \geq 8$) does the optimal portfolio include T-bills.

Why does the introduction of housing push the household into a situation in which the non-negativity constraint binds over a wider range of risk aversion? For the first three age cohorts (that is, households aged 50 and under), the value of the house exceeds net wealth ($h > 1$), so that the household must borrow. Risk-tolerant investors exploit the opportunity to leverage their asset portfolio and take out the maximum mortgage equal to 100 percent of the value of the house while more risk-averse households avoid being fully leveraged, but all of the portfolios associated with values of h exceeding unity involve substantial mortgage indebtedness. If allowed, diversifying the portfolio of financial assets by taking a negative position in T-bills and reducing mortgage indebtedness would both increase the portfolio's expected return and reduce its variance, since the interest rate on T-bills is, on average, slightly lower than the mortgage rate. For households in the last three age cohorts, owning the house without holding any mortgage is feasible ($h < 1$). However, except for the highly risk averse ($A = 10$), these households voluntarily hold mortgages in order to hold larger amounts of the high-return financial assets (stocks and bonds) than would be possible if they owned their house outright. Like the younger, "mandatory" mortgage holders, these "voluntary" mortgage holders would be better off if some of their borrowing were in the form of a negative position in T-bills. Since both of these groups would like to borrow at the T-bill rate, the nonnegativity constraint is binding and the optimal holding of T-bills is zero.

To assess the empirical validity of this cross-sectional implication of the model—that the inclusion of housing pushes the household more

firmly onto a binding nonnegativity constraint for T-bills—we reexamine the "asset allocation puzzle" as posed by Canner et al. (1997). In essence, the asset allocation puzzle refers to the following contradiction between theory and practice: according to the mutual fund separation theorem, all investors should hold an identical portfolio of risky assets regardless of their degree of risk aversion; differential preferences toward risk are accommodated by varying the relative amounts invested in the risky portfolio and the riskless asset, which implies that the ratio of bonds to stocks (two risky assets) should be invariant with respect to the overall level of risk of the portfolio, as measured by the percentage of the entire portfolio held in the form of stocks. In contrast, financial advisors in the popular press systematically advise that more risk-averse investors hold higher ratios of bonds to stocks than risk-tolerant investors.

Instead of using data on actual household portfolios, Canner et al. use as "data" the portfolios recommended by four sources of investment advice: a book by Jane Bryant Quinn, the Fidelity investment newsletter, a book promoted by Merrill Lynch, and an article in the *New York Times*.¹⁴ Table 4 reproduces Table 1 of Canner et al. (1997). While the sources differ in their recommended portfolios, all four advise a ratio of bonds to stocks which increases with risk aversion, and is therefore inversely related to the fraction of the portfolio in stocks.

Parallel to Figure 1 in Canner et al., Figure 2 plots (with circle symbols) the bond-to-stock ratio against the percentage of financial assets in stocks for the 12 recommended portfolios.¹⁵ The two dashed lines in Figure 2 represent the predicted relationship between the bond-to-stock ratio and the percentage of financial assets in stocks, based on the standard CAPM assumptions (i.e., a riskless asset exists, short sales are permitted, and housing is excluded) and on a variant of the CAPM in which there is no risk-

¹⁴ The sources are: Jane Bryant Quinn (1991), Larry Mark (1993), Don Underwood and Paul B. Brown (1993), and Mary Rowland (1994).

¹⁵ There are only ten distinct portfolios because Fidelity and Jane Bryant Quinn recommend the identical portfolio allocation for conservative investors, and the *New York Times* "conservative" portfolio coincides with the Fidelity "moderate" portfolio.

TABLE 4—ASSET ALLOCATIONS RECOMMENDED BY FINANCIAL ADVISORS

Advisor and Investor Type	Percent of Portfolio			Ratio of Bonds to Stocks
	Cash	Bonds	Stocks	
A. Fidelity				
Conservative	50	30	20	1.50
Moderate	20	40	40	1.00
Aggressive	5	30	65	0.46
B. Merrill Lynch				
Conservative	20	35	45	0.78
Moderate	5	40	55	0.73
Aggressive	5	20	75	0.27
C. Jane Bryant Quinn				
Conservative	50	30	20	1.50
Moderate	10	40	50	0.80
Aggressive	0	0	100	0.00
D. <i>The New York Times</i>				
Conservative	20	40	40	1.00
Moderate	10	30	60	0.50
Aggressive	0	20	80	0.25

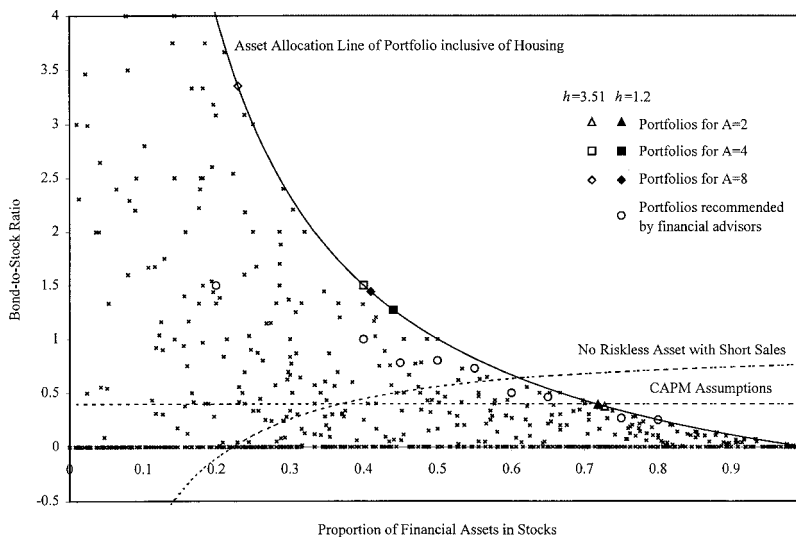


FIGURE 2. PLOT OF BOND-TO-STOCK RATIO AGAINST THE PROPORTION OF FINANCIAL ASSETS IN STOCKS

less asset (while short sales are permitted, and housing is excluded). These elements of Figure 2 illustrate the “asset allocation puzzle” as stated by Canner et al.: when housing is excluded from the model, the predicted relationship between the bond-to-stock ratio and the percentage of financial assets in stocks is exactly horizontal (if a riskless asset exists) or positively sloped (in the absence of a riskless asset), while the scatterplot representing the fi-

nancial advisors’ recommendations is negatively sloped.¹⁶

Canner et al. examine a number of potential

¹⁶ Canner et al. (1997) calculated optimal portfolios based on pretax real returns, while we use after-tax real returns. Because of the difference in tax treatment, the constant bond-to-stock ratio consistent with the “CAPM assumptions” is equal to 0.40 in our calculations, as compared to 0.33 in Canner et al. (1997).

explanations for the puzzle, but find that none of “the usual suspects” can successfully resolve the puzzle. The authors conclude¹⁷

[T]he failure of [portfolio] advice, as opposed to behavior, strikes us as particularly surprising. Financial decisions are widely viewed to be difficult to make, and advisors are experts who are trying to help people optimize. That the advice being offered does not match economic theory suggests that our understanding of investor objectives (as opposed to their ability to reach those objectives) is deficient.

Once we introduce housing into the portfolio, for values of h greater than 1.2, the predicted relationship between the bond-to-stock ratio and the percentage of financial assets in stocks is given by the negatively sloped solid curve in Figure 2.¹⁸ For a wide range of values of risk aversion (even for $A > 10$), the model predicts that the nonnegativity constraint on T-bills is binding as long as h exceeds 1.2. Since the household is at a corner with zero T-bills, the percentage of financial assets in bonds and the percentage of financial assets in stocks must sum to unity in this region. Thus, if we denote the percentage of financial assets in stocks as s , the prediction of the model is simply a plot of $(1 - s)/s$ against s , a negatively sloped curve. As noted earlier, the introduction of housing does not universally imply that the nonnegativity constraint on T-bills is binding. For households with both low levels of h and high levels of A (for example, $h = 0.65$ and $A = 8$) the optimal portfolio does contain T-bills and therefore would be represented by a point to the southwest of the curve which represents the plot of $(1 - s)/s$ against s . Nevertheless, it is evident from Table 2 that for a wide range of combinations of h and A , the binding nonnega-

tivity constraint on T-bills implies that the optimal portfolio will lie on this curve, labeled “asset allocation with housing.”

The mean-variance efficiency framework assumes zero transactions costs in buying and selling financial assets, and does not recognize a distinction between financial assets which can be used as a medium of exchange, such as checking accounts and currency, as opposed to nonmonetary assets. The implication of the model is that, for most homeowners, the optimal portfolio consists entirely of bonds and stocks; no cash is held for the purpose of achieving mean-variance efficiency. Before confronting the portfolio implications of the model with the data, however, it is necessary to consider the implications of transactions costs. In the absence of transactions costs, the household would hold zero cash and pay for consumption goods by selling bonds or stocks. However, in the presence of transactions costs of buying and selling bonds and stocks, the household will hold some cash balances, even though the non-zero holdings of cash move the household to the interior of the mean-variance efficient frontier (as computed ignoring transactions costs). Since the “asset allocation line with housing” represents the maximum ratio of bonds to stocks for a given percentage of financial assets in stocks, a portfolio which includes a strictly positive amount of cash will be represented by a point below the curve.

In addition to the “data” which represents the portfolios recommended by the financial advisors, Figure 2 also plots, with small x’s, data points representing the portfolios of PSID homeowners in 1989.¹⁹ We interpret the data—both the scatter of household portfolios and the financial advisors’ recommended portfolios—as deviating from the prediction of the frictionless model because cash is held for transactions purposes. The inescapable fact that households generally hold strictly positive balances of cash does not contradict the model’s implication that most households face a binding nonnegativity constraint on T-bills, or “cash.” If holding negative positions in checking accounts, currency, or T-bills were feasible, the household could

¹⁷ Canner et al. (1997 p. 189).

¹⁸ The target audience of the financial advisors is presumably the subset of the overall population with the wealth and motivation to hold diversified portfolios of financial assets. The rate of homeownership among this group is even higher than in the general population. For example, of PSID households with strictly positive stock holdings in 1989, 93 percent of those in the 41–70 age range were homeowners. Of stockholders aged 31–40, 82 percent were homeowners, and of those aged 71 and older, 88 percent were homeowners.

¹⁹ The asset labeled “T-bills” in the predicted portfolios is interpreted, in terms of the PSID data, as the sum of checking, savings, and money market accounts, as well as T-bills.

hold a positive balance of the cash asset to use as a medium of exchange and simultaneously borrow by holding negative balances of checking accounts or T-bills and thus accommodate a transactions demand for cash while holding a negative net position in the cash asset to achieve mean-variance efficiency. Since borrowing at the currency or T-bill rate is not possible, the household accepts some costs in terms of mean-variance inefficiency in order to hold transactions balances. The fact that the household would like to simultaneously take a negative position in the cash asset to achieve mean-variance efficiency, but is unable to do so, implies that the nonnegativity constraint is nevertheless binding.

IV. Conclusions

Simply by virtue of its magnitude, housing plays an important role in both the consumption bundle and the asset portfolio of the household. This role is complicated, and made more interesting, by the fact that most consumers choose to hold a single level of residential real estate to satisfy both the consumption demand for housing services as well as the portfolio demand for housing, rather than use rental markets to disentangle the two aspects of the problem. To the extent that the household's holding of real estate is determined at least in part by its consumption demand for housing services, the consumption demand for housing places a constraint on the portfolio problem. Using PSID data and the Case-Shiller repeat sales transactions price data, the paper estimates the risk and return to financial assets and residential real estate, and calculates the optimal portfolios and constrained mean-variance efficient frontiers for various values of the housing constraint.

The model implies that households which are *ex ante* identical, in the sense that they have identical preferences toward risk and identical perceptions of the risk and return to different assets, will nevertheless hold quite different portfolios of financial assets because each household is optimizing their portfolio subject to a constraint on housing, and this constraint varies across households. Further, since the ratio of housing to net worth declines over the life cycle, the housing constraint generally induces a life-cycle pattern in the portfolio shares of stocks and bonds. For example, young house-

holds, which typically have large holdings of real estate relative to their net worth, are highly leveraged and therefore forced into a situation of high portfolio risk (and return). As a result, these young households respond to the housing constraint by using their net worth to either pay down their mortgage or buy bonds instead of buying stocks. In comparison, ownership of stocks is more attractive to older households that have accumulated greater wealth and therefore reduced their ratio of housing to net worth.

Compared to the mean-variance frontier which corresponds to portfolios consisting solely of financial assets such as stocks, bonds, and T-bills, the introduction of housing and mortgages alters the risk and return trade-off in a direction which pushes most households onto a binding nonnegativity constraint on T-bills. Since households are at a corner with respect to the riskless, or low-risk, asset, varying preferences toward risk cannot be accommodated by altering the percentage of the portfolio in the riskless asset, and are reflected instead in variation in the composition of the risky portfolio. This observation—that the introduction of housing pushes the household onto a binding nonnegativity constraint on the riskless asset—appears to be useful in explaining both household portfolio data and data points which represent the portfolio recommendations of financial advisors in the popular press.

The analysis in the paper is entirely partial equilibrium, and assumes that the expected return vector and covariance matrix are time invariant. Nevertheless, it is fun to speculate on the consequences of the housing constraint in a dynamic, general-equilibrium setting. If one accepts the basic premise that the household's optimal portfolio is constrained by the housing ratio, and this ratio tends to be related to age, then demographic bulges such as the baby boom, birth dearth, and baby boom echo may have important effects on asset prices. The effect of the baby boom generation on the expected return to housing has been demonstrated in the literature (Joyce Manchester, 1989; N. Gregory Mankiw and David N. Weil, 1989). Similarly, demographic effects have been invoked as a possible explanation of the recent stock market boom. However, the treatment of demographic effects on the stock market in the popular media generally focuses on the idea that the baby boomers, finally recognizing that

retirement is on the horizon, have comparatively high saving rates. In this view, the aggregate flow of saving is unusually high because the baby boom cohort has a high saving rate and constitutes a large fraction of the overall population. Our analysis focuses on portfolio shares rather than saving rates, and shows that as the housing constraint is gradually relaxed over the life cycle, the aging baby boomers should shift their portfolio composition away from bonds and toward stocks. Thus even in the absence of high saving rates, the baby boom generation could have a systematic effect on asset prices.

DATA APPENDIX

Calculation of Housing Return from PSID.—Of the 6,771 households in the Survey Research Center (SRC) (nonpoverty) sample in 1992, households that never owned a home between 1968 and 1992 were excluded, reducing the sample to 5,846. Households for which the value of the house or remaining mortgage principal have ever been assigned by PSID staff between 1968 and 1992 were also excluded, reducing the sample to 1,817. Households that enter the sample as a split-off of another family unit are included, but data for the years during which the split-off household resided with the original family unit are treated as missing.

Inspection of the PSID house value data revealed 45 cases in which the house value changed by a factor of ten from one year to the next even though the family unit did not move. In these cases, examination of the series of reported house values in the context of the rest of the record led us to conclude that the observation had been incorrectly recorded. These observations were corrected by either adding or dropping a zero, and included in the sample. A list of the observations for which the house value was corrected is available from the authors on request.

If the family unit owns the house they reside in, the house value is elicited by the question “Could you tell me what the present value of your house is?—I mean about how much would it bring if you sold it today?” The property tax amount is also a direct survey question except in 1978, 1988, and 1989. For 1978, we impute property taxes by the method suggested in the PSID Wave VI Codebook (Survey Research Center, 1973). Property tax rates are assigned

TABLE A1—PROPERTY TAX RATE

Distance from Nearest City of 50,000 or more	New England States (percent)	All Other States (percent)
0–5 miles	2.5	2.0
5–49 miles	2.0	1.5
50 or more miles	1.5	1.0

Source: Survey Research Center (1973 p. 140).

according to Table A1, based on the location of the house. For 1988 and 1989, we lack data on the distance from the nearest city over 50,000 and use instead the size of the largest city in the county as an indicator. For 1988–1989, the property tax rate was assumed to be 1.5 percent (2.0 percent for New England states) if the largest city in the county is greater than 50,000, and 1.0 percent (1.5 percent for New England states) otherwise.

In calculating the housing return, the observation was set to missing if during the interval over which the return was calculated, the family moved or there was a change of head due to marriage of a female head of household to a nonsample male. Since the number of homeowners in the PSID changes every year, we construct 24 vectors of returns and covariance matrices (one for each year) and then weight them equally to arrive at a single vector of returns and covariance matrix for the 1968–1992 sample period.

Center for Research on Security Prices (CRSP) Data.—Returns for stocks, bonds, and bills are calculated from monthly data series provided by Ibbotson and Associates. For additional information, see *CRSP Indices File Guide* (Center for Research on Security Prices, 1996).

Treasury Bills (USTRET)—U.S. Treasury Bills total return.

Treasury Bond (GBTRET)—U.S. Treasury Bonds total return.

Stocks Capital Appreciation (CSCRET)—S&P 500 stocks capital appreciation. The capital appreciation component of the common stock return is taken as the change in the S&P 500 as reported in the *Wall Street Journal* over 1977–1992, and in Standard and Poor’s Trade Securities Statistics from 1926–1976.

TABLE A2—RETURN AND COVARIANCE MATRIX—PSID DATA WITH TIME-VARYING FEDERAL MARGINAL TAX RATES FOR 1969–1992

	T-Bills	Bonds	Stocks	Mortgage	House	
Mean return (arithmetic)	-0.0045	0.0051	0.0820	-0.0019	0.0662	
Standard deviation	0.0434	0.0814	0.2415	0.0335	0.1424	
Covariance matrix						
	T-bills	0.0018854				
	Bonds	0.0024652	0.0066290			
	Stocks	0.0002202	0.0038510	0.0583161		
	Mortgage	0.0006995	0.0023519	0.0024675	0.0011208	
	House	-0.00012	-0.000072	-0.000017	-0.00001	0.0202878
Correlation matrix						
	T-bills	1.00000				
	Bonds	0.69731	1.00000			
	Stocks	0.02100	0.19586	1.00000		
	Mortgage	0.84436	0.66799	0.45706	1.00000	
	House	-0.03413	-0.00483	-0.00742	-0.00056	1.00000

Stocks Income Return (CSIRET)—S&P 500 stocks income return. For 1977–1992, income is calculated as realized dividends. Dividends are accumulated over the month and then invested on the last trading day of the month in the S&P 500 at the day's closing level. For 1926–1976, quarterly dividends are extracted from rolling yearly dividends reported quarterly in Standard and Poor's Trade and Securities Statistics, then allocated to months within each quarter, assuming payments were made on the same month and day as they were in 1974.

CPI (CPIRET)—Consumer Price Index for All Urban Consumers, not seasonally adjusted (CPI-U NSA).

Mortgage Rates.—The data for mortgage rates is the variable "Conventional Home Mortgage Rate—Loans Closed (National Average for All Major Lenders: FYMCLE)" from the Citibase Monthly data tape. The annual average (April to March) of the monthly mortgage rate is converted to an after-tax rate using the marginal tax rate of 33 percent, and then deflating by the inflation rate.

In Table A2, we report an alternative version of the expected returns and covariance matrix of returns, in which historical variation over time of the federal marginal tax rate is included. That is, for each year, we calculated the marginal federal tax rate of the PSID household that ranked in the 75th percentile for income, and used the year-specific marginal federal tax rate

to calculate after-tax real returns. A comparison of the estimates with the corresponding estimates in Table 1A of the paper indicates that the alternative method of incorporating taxes makes very little difference.

REFERENCES

- Ameriks, John and Zeldes, Stephen P.** "How Do Household Portfolio Shares Vary with Age?" Unpublished manuscript, Columbia University, 2000.
- Bertaut, Carol C.** "Stockholding Behavior of U.S. Households: Evidence from the 1983–89 Survey of Consumer Finances." International Finance Discussion Papers No. 558, Board of Governors of the Federal Reserve System, Washington, DC, July 1996.
- Brueckner, Jan K.** "Consumption and Investment Motives and the Portfolio Choices of Homeowners." *Journal of Real Estate Finance and Economics*, October 1997, 15(2), pp. 159–80.
- Canner, Niko; Mankiw, N. Gregory and Weil, David N.** "An Asset Allocation Puzzle." *American Economic Review*, March 1997, 87(1), pp. 181–91.
- Case, Karl E. and Shiller, Robert J.** "The Efficiency of the Market for Single-Family Homes." *American Economic Review*, March 1989, 79(1), pp. 125–37.
- Center for Research in Security Prices.** *CRSP indices file guide*, 1996.
- Cocco, João F.** "Owner-Occupied Housing,

- Permanent Income, and Portfolio Choice." Unpublished manuscript, Harvard University, 1998.
- Fratantoni, Michael.** "Housing Wealth, Precautionary Saving and the Equity Premium." Unpublished manuscript, Johns Hopkins University, 1996.
- Goetzmann, William N.** "The Single Family Home in the Investment Portfolio." *Journal of Real Estate Finance and Economics*, May 1993, 6(3), pp. 201–22.
- Goetzmann, William N. and Ibbotson, Roger.** "The Performance of Real Estate as an Asset Class." *Journal of Applied Corporate Finance*, Spring 1990, 3(1), pp. 65–76.
- Goodman, John L., Jr. and Ittner, John B.** "The Accuracy of Home Owners' Estimates of House Value." Economic Activity Section Working Paper No. 131, Board of Governors of the Federal Reserve System, Washington, DC, 1993.
- Grossman, Sanford J. and Laroque, Guy.** "Asset Pricing and Optimal Portfolio Choice in the Presence of Illiquid Durable Consumption Goods." *Econometrica*, January 1990, 58(1), pp. 25–51.
- Heaton, John and Lucas, Deborah.** "Portfolio Choice in the Presence of Background Risk." *Economic Journal*, January 2000, 110(460), pp. 1–26.
- Henderson, J. Vernon and Ioannides, Yannis M.** "A Model of Housing Tenure Choice." *American Economic Review*, March 1983, 73(1), pp. 98–113.
- Ioannides, Yannis M.** "Dynamics of the Composition of Household Asset Portfolios and the Life Cycle." *Applied Financial Economics*, September 1992, 2(3), pp. 145–59.
- King, Mervyn A. and Leape, Jonathan I.** "Asset Accumulation, Information, and the Life Cycle." National Bureau of Economic Research (Cambridge, MA) Working Paper No. 2392, September 1987.
- Manchester, Joyce.** "The Baby Boom, Housing and Loanable Funds." *Canadian Journal of Economics*, February 1989, 22(1), pp. 128–49.
- Mankiw, N. Gregory and Weil, David N.** "The Baby Boom, the Baby Bust, and the Housing Market." *Regional Science and Urban Economics*, May 1989, 19(2), pp. 235–58.
- Mark, Larry.** "Asset Allocation: Finding the Right Mix." *Fidelity Focus: The Magazine for Fidelity Investors*, Winter 1993, p. 11.
- Poterba, James M. and Samwick, Andrew A.** "Household Portfolio Allocation over the Life Cycle." National Bureau of Economic Research (Cambridge, MA) Working Paper No. 6185, 1997.
- Quinn, Jane Bryant.** *Making the most of your money*. New York: Simon and Schuster, 1991.
- Ross, Stephen A. and Zisler, Randall C.** "Risk and Return in Real Estate." *Journal of Real Estate Finance and Economics*, June 1991, 4(2), pp. 175–90.
- Rowland, Mary.** "Seven Steps to Handling an Inheritance." *The New York Times*, February 5, 1994, p. 34.
- Skinner, Jonathan.** "Housing and Saving in the United States," in Yukio Noguchi and James M. Poterba, eds., *Housing markets in the United States and Japan*. Chicago: University of Chicago Press, 1994, pp. 191–213.
- Survey Research Center.** *A panel study of income dynamics, procedures and tape codes: 1973 interviewing year, wave VI, a supplement*. Ann Arbor, MI: Institute for Social Research, University of Michigan, 1973.
- Underwood, Don and Brown, Paul B.** *Grow rich slowly: The Merrill Lynch guide to retirement planning*. New York: Viking, 1993.
- Yamashita, Takashi.** "Owner-Occupied Housing and Investment in Stocks: An Empirical Test." Unpublished manuscript, University of Nevada, Las Vegas, 2001.