The Effect of Housing on Portfolio Choice*

Raj Chetty
Harvard and NBER

Adam Szeidl
UC-Berkeley and NBER

May 2010

Abstract

A large theoretical literature predicts that housing has substantial effects on financial markets, but empirical evidence on these effects remains limited. We estimate the causal effect of changes in mortgages and home equity on portfolio allocations using two empirical strategies. First, we use two instruments – average house prices in an individual's state in the current year and in the year he purchased his home – to generate cross-sectional variation in home equity and mortgages that is plausibly orthogonal to unobserved determinants of portfolios. Second, we use panel data to study how portfolio allocations change when individuals buy houses. Both empirical strategies show that housing reduces the amount households invest in stocks substantially: a $10,000 increase in mortgage debt (holding fixed total wealth) reduces the stock share of liquid wealth by approximately 6%. Auxiliary evidence suggests that housing induces individuals to hold more conservative portfolios primarily because of a "consumption commitment" effect rather than exposure to house price risk.

*E-mail addresses: chetty@fas.harvard.edu and szeidl@econ.berkeley.edu. Thanks to Thomas Davidoff, Edward Glaeser, Stephen Shore, and numerous seminar participants for helpful comments. Gregory Bruich, Keli Liu, James Mahon, Juan Carlos Suarez Serrato, and Philippe Wingender provided outstanding research assistance. We are grateful for funding from National Science Foundation Grant SES 0522073.

1 Introduction

How does homeownership affect households' financial investment decisions? The interaction between housing and financial markets has attracted attention because of its importance for understanding macroeconomic fluctuations and asset pricing. Theoretical studies have shown that housing affects optimal portfolio allocations through two channels because it is both a consumption good and an illiquid asset. First, home ownership increases a household's exposure to risk (Brueckner 1997, Flavin and Yamashita 2002). Under plausible assumptions, this can induce a household to hold a more conservative financial portfolio. Second, adjustment costs in housing effectively amplify risk aversion because they force households to concentrate fluctuations in wealth of a subset of goods. This "consumption commitment" effect can further reduce the optimal share of stocks in liquid wealth (Grossman and Laroque 1990, Fratantoni 2001, Chetty and Szeidl 2007). Together, these effects can have large quantitative effects on portfolios. Simulations by Cocco (2005) show that housing can reduce stock market participation rates from 76% to 33% in a calibrated life-cycle model.

The goal of this paper is to test these theoretical predictions by estimating the effect of owning a more expensive house on portfolio choice empirically. Previous empirical work has documented the cross-sectional correlation between property values and portfolio choice using OLS regressions (Fratantoni 1998, Heaton and Lucas 2000, Yamashita 2003, Cocco 2005).[1] Because both portfolios and housing are endogenous choices that are affected by unobserved factors such as background risk (Campbell and Cocco 2003, Cocco 2005, Davidoff 2009), these studies may not identify the causal effect of housing on portfolios.

In this paper, we use two empirical strategies to address this endogeneity problem: a cross-sectional instrumental variables (IV) approach and a panel analysis of changes in portfolios around housing purchases. We motivate our empirical analysis using a simple, analytically tractable model of portfolio choice that incorporates both the commitment and price risk effects of housing. In our model, agents maintain a buffer stock of liquid wealth while holding a home mortgage in order to smooth consumption in the presence of borrowing

[1] Yamashita (2003) also reports some two-stage-least-squares estimates using age, family size, home tenure, and aggregate housing returns as instruments. Unfortunately, standard models would generate direct relationships between all of these variables and portfolio choice, independent of the housing channel. Our study builds on Yamashita's work by using instruments that are more likely to satisfy the exclusion restriction and by separating the effects of mortgage debt and home equity wealth.

costs. Housing induces agents to hold this buffer stock in more conservative assets. A key implication of the model is that it is important to distinguish changes in mortgage debt from changes in home equity wealth to characterize the effects of housing on portfolios. Increases in property value should reduce the stock share of liquid wealth through the combination of the two housing mechanisms, while home equity increases stockholdings through a wealth effect.[2]

Our first empirical strategy instruments for property values and home equity using current and year-of-purchase home prices in the individual's state, calculated using repeat-sales indices. The current house price index is naturally a strong predictor of property values. However, the current house price also creates variation in a household's wealth: increases in house prices increase home equity wealth. To isolate the causal effect of a more expensive house while holding wealth fixed, we exploit the second instrument – the average house price at the time of purchase. Individuals who purchase houses at a point when prices are high tend to have less home equity and a larger mortgage. Critically, we control for aggregate shocks and cross-sectional differences across housing markets by including state and year fixed effects, thereby exploiting only differential within-state variation for identification. With these fixed effects, the two state house price index instruments generate variation in property value and home equity wealth that is orthogonal to several unobserved determinants of financial portfolios.

We implement this cross-sectional IV strategy using microdata on housing and portfolios for 69,130 households from the Survey of Income and Program Participation (SIPP) panels spanning 1990 to 2004. We use two-stage-least-squares specifications to estimate the effect of property value and home equity on the share of liquid wealth that a household holds in stocks. We find that housing has a large effect on financial portfolios: a $10,000 increase in property value (holding fixed home equity wealth) causes a reduction in the stock share of liquid wealth of approximately 6%. This estimate is stable and statistically significant with $p < 0.05$ across a broad range of specifications. In contrast, a $10,000 increase in home equity (holding fixed total property value) increases the stock share of liquid wealth by 6.5% through a wealth effect. Hence, it is critical to distinguish changes in mortgage debt from the wealth effect of

[2] Because property value is the sum of mortgage debt and home equity, changes in property value holding fixed home equity are equivalent to changes in mortgage debt.

home equity to characterize the effects of housing on portfolios. The portfolio changes we document are driven by a combination of extensive and intensive-margin responses. Larger mortgages induce reductions in both the probability of stock market participation and the amount of stocks held conditional on participation.

Our second empirical strategy uses panel data to examine changes in portfolios from the year before to the year after home purchase. In particular, we test whether individuals who buy a larger house reduce their stock share of liquid wealth more than those who buy smaller houses. This strategy mitigates the endogeneity of housing choice in the cross-section by permitting household fixed effects. Since the amount an individual spends on a house may be endogenous to changes in risk preferences or wealth, we also estimate models where we instrument for the change in property value using the state-level house price index at the time of home purchase. This panel strategy complements the cross-sectional IV approach in two ways. First, it provides evidence that households actively change the composition of their financial portfolios depending upon the amount they invest in a house. Second, it relaxes the exclusion restriction underlying the first identification strategy – namely that current house prices and house prices at the time of purchase do not affect portfolios directly. The second strategy effectively shows that house price indices have no effect on portfolio shares *prior* to home purchase, supporting the exclusion restriction underlying the cross-sectional analysis.

Because the SIPP is a short panel, we observe portfolio shares both before and after home purchase for only 2,784 households. We implement the second empirical strategy by examining the change in stock shares of liquid wealth from the year before to the year after home purchase for this subset of households. Using an OLS first-differences specification, we find that a household that buys a $10,000 more expensive house reduces stockholding by a statistically significant 3% of liquid wealth. Instrumenting for the price of the house with the state home price index at the time of purchase, we find that a $10,000 increase in the price of the house leads to a 9% reduction in the stock share of liquid wealth. Stock shares fall because households primarily sell stocks to finance larger downpayments: OLS estimates imply that a $10,000 increase in property value leads to a $3,300 drop in liquid wealth, of which $2,650 comes from the sale of stocks and $240 comes from the sale of safe assets.

Combining the empirical estimates from the two strategies, we conclude that a $10,000 increase in mortgage debt reduces stock holdings by approximately 6% on average. This estimate implies that if all mortgage debt were forgiven, households would increase the amount of money invested in stocks by 30%, ignoring general equilibrium effects on asset prices.

The results described thus far do not tell us whether the effect of housing on portfolios is driven by the home price risk or commitment channels. To distinguish between these two channels, we examine how the effect of housing on portfolios varies across markets with different levels of house price risk and adjustment costs. The home price risk hypothesis predicts that the effect of property value on stock shares will be more negative in markets with more volatile house prices. The commitment channel predicts that the effect should be larger in more illiquid housing markets. We proxy for house price risk using historical house price volatility by state and for adjustment costs using average home tenures by state. The estimated effects of property value on stock shares are very similar in low- and high-risk states, but are significantly higher for individuals living in states with high home tenures (i.e. those with higher adjustment costs). These results suggest that the commitment effect is the more important driver of the link between housing and portfolios. This is consistent with Sinai and Souleles' (2005) finding that most individuals are well hedged against house price risk because they do not move across markets frequently. However, these auxiliary tests must be interpreted only as suggestive correlations, because they rely on cross-sectional variation in house price risk and adjustment costs.

Our estimates of the effect of housing on portfolios – an elasticity of -0.3 of stock shares with respect to mortgage debt and 0.44 of stock shares with respect to home equity wealth – are substantially larger in magnitude than those of previous cross-sectional studies. The results of previous studies are sensitive to controls and sample specification partly because they do not directly address the endogeneity of housing. For example, Fratantoni (1998) finds an elasticity of stock share with respect to mortgage debt of -0.15. In contrast, Heaton and Lucas (2000) and Cocco (2005) show that once property value is included as a covariate, mortgages are *positively* associated with the stock share in cross-sectional OLS regressions. Yamashita (2003) finds an elasticity of stock share with respect to property value of approximately -0.1 in a specification that does not include mortgage debt. Consistent with prior studies, we show that cross-sectional OLS estimates in our data are sometimes wrong-signed and are very sensitive to covariates. In contrast, our instrumental-variable estimates are not very sensitive to covariates because they isolate variation that is orthogonal to most household-level determinants of portfolios. Moreover, our analysis shows that it is critical to distinguish home equity wealth from mortgage debt to understand the effect of housing on portfolios. The mixed results of prior studies are also driven by the way in which they treat this distinction.

The reduced-form empirical results here point to potentially important linkages between housing and asset markets. For example, our results suggest that recent increases in mortgage debt relative to liquid wealth in the U.S. may have induced households to withdraw funds from the stock market. This shift in demand for risky assets could have further precipitated the sharp decline in asset prices. A full quantitative analysis of the effect of housing market shocks on asset markets would require a general equilibrium model of financial and real estate markets. The elasticities we report here can be used when calibrating such a model, and the model we present below could provide a tractable starting point for such work.

The remainder of the paper is organized as follows. The next section presents a portfolio choice model and derives an estimating equation for the empirical analysis. Section 3 describes the data and section 4 presents the identification strategies and results. Section 5 concludes.

2 Model and Estimating Equation

This section develops a model that incorporates house price risk and moving costs in a simple portfolio choice framework. The objective of this highly stylized model is to highlight the channels through which housing affects portfolios and guide the specification of an estimating equation.

Consider a discrete-time environment with periods $t = 0, 1, ..., T$. The agent consumes only in periods $t = 1, ..., T$ and maximizes total consumption utility

$$\sum_{t=1}^{T} \frac{c_t^{1-\gamma}}{1-\gamma}$$

where we abstract from discounting and assume that the coefficient of relative risk aversion $\gamma > 1$. Following Cocco (2005), let c_t denote a Cobb-Douglas composite of food (f) and housing (h) consumption. Housing h is itself a composite of two types of durable goods, commitments (x) and adjustables (a), so that

$$c = f^{\mu}h^{1-\mu} \text{ and } h = \left(\frac{x}{\theta}\right)^{\theta}\left(\frac{a}{1-\theta}\right)^{1-\theta}. \tag{1}$$

The parameter θ governs the adjustability of housing consumption. The committed component of housing (x) can never be adjusted from its prior level, while component a is freely adjustable at all dates $t \geq 1$. Hence, if $\theta = 0$, all housing is instantaneously adjustable; if $\theta = 1$, housing consumption is completely fixed.[3] At $t = 0$, the household is endowed with L_0 dollars of liquid wealth (cash); an initial house of size h_0, of which $x_0 = \theta h_0$ is unadjustable and $a_0 = (1-\theta)\,h_0$ is adjustable; and a mortgage of face value M_0. All these variables are exogenously fixed – the agent makes choices over housing and food only starting $t = 1$. Moreover, the commitment portion of housing is fixed for the lifetime of the household: $x_t = x_0$ for all t. This stark method of modelling consumption commitments allows us to obtain a closed form portfolio choice equation.

The agent can invest in two assets in each period $t = 0, 1, ..., T$: stocks and a risk free bond. We normalize the risk free borrowing and lending rate to zero. All uncertainty is realized between $t = 0$ and $t = 1$. Hence, stocks earn a zero return after $t = 1$ and the only non-trivial portfolio choice problem is at $t = 0$. Let R denote the net return on stocks between $t = 0$ and $t = 1$, and assume that $\log(1 + R)$ is normally distributed with variance σ^2. The agent has a riskless labor income stream of y_t; let $Y = \sum_{t=1}^{T} y_t$ denote the present value of labor income.

In addition to the stock market, the agent faces risk from fluctuations in the price of housing between $t = 0$ and $t = 1$.[4] We assume that the price of adjustable and commitment housing (both durable goods) are equal, and denote this price in period t by p_t. We further simplify the model by assuming that housing and the stock market are perfectly correlated, so that the return on housing (p_1/p_0) from period 0 to 1 is replicated by a financial portfolio that invests some share κ in stocks and $1-\kappa$ in bonds. With this specification, κ can be interpreted as a measure of home price risk: greater comovement with the stock market

[3] We normalize the denominator of the expression for h by $\theta^{\theta}(1-\theta)^{1-\theta}$ to ensure that a dollar of housing expenditure generates the same amount of housing consumption for all values of θ.

[4] It would be straightforward to introduce background risk in labor income that is perfectly correlated with the stock market.

makes housing riskier. When $t \geq 1$, the per period rental price of both types of housing is p_1/T.

Denoting the share of liquid wealth invested in stocks at $t = 0$ by α, the agent's maximization problem is thus

$$\max_{\alpha, f_1, \ldots, f_T, a_1, \ldots, a_T} E \sum_{t=1}^{T} \frac{c_t^{1-\gamma}}{1-\gamma}$$

subject to (1) and the budget constraint

$$(1 + \alpha R) L_0 + Y + (p_1 h_0 - M_0) = \sum_{t=1}^{T} (f_t + a_t \cdot p_1/T + x_0 \cdot p_1/T).$$

The budget constraint requires that the sum of financial wealth, labor income, and home equity, all measured at $t = 1$ (the left hand side), equals the sum of food consumption, adjustable housing consumption and committed housing consumption (the right hand side). The committed component of housing consumption, appearing on the right hand side, effectively reduces the disposable income available for food consumption as income fluctuates. Home price risk affects the agent's budget constraint in two ways: on the left hand side, it introduces background wealth risk in that the value of the initial property $p_1 h_0$ fluctuates; on the right hand side, it introduces risk in the relative price of housing and food (p_1/T) in periods $t \geq 1$.

The following proposition characterizes the agent's optimal portfolio choice at $t = 0$ (α) as a function of initial liquid wealth (L_0), home equity ($p_0 h_0 - M_0$), property value ($p_0 h_0$), and total labor income (Y).

Proposition 1 *The optimal share of stocks out of liquid wealth at $t = 0$ is, to a log-linear approximation,*

$$\alpha^* = C_1 \cdot \frac{\textit{liquid wealth} + \textit{labor income} + \textit{home equity}}{\textit{liquid wealth}} - [\theta C_1 + \kappa C_2] \cdot \frac{\textit{property value}}{\textit{liquid wealth}} \quad (2)$$

where C_1 and C_2 are non-negative constants:

$$C_1 = \frac{ER - (1-\gamma)\,\mu_a \cdot \kappa\sigma^2}{\sigma^2 \left[1 - (\mu + \mu_a)(1-\gamma)\right]} \quad \textit{and} \quad C_2 = 1 - \theta$$

and $\mu_a = (1-\mu)(1-\theta)$ is the Cobb-Douglas share of adjustable housing consumption.

To understand the proposition, first consider the special case when $\theta = \kappa = 0$, i.e., when all of housing is adjustable and there is no home-price risk. In that case the final term in (2) is zero; stockholdings are a constant share of total wealth, as in the standard power utility model (Campbell and Viceira, 2002). With either home price risk ($\kappa > 0$) or commitments ($\theta > 0$), stockholding is negatively related to property value, conditional on total wealth. In the presence of home price risk, a larger house translates into larger background risk, resulting in lower stockholdings. In the presence of commitments, a larger home means that more wealth is "tied up" in fixed housing payments; hence a given loss in wealth results in a larger proportional drop in food consumption, driving up risk aversion and reducing stockholdings (Chetty and Szeidl 2007a).

The degree of home price risk and commitments affect the coefficient on property value in (2) in intuitive ways. An increase in home price risk (higher κ) amplifies the effect of property value on portfolios because every additional dollar of housing leads to greater exposure to price risk.[5] When housing is less adjustable (higher θ), the property value coefficient changes in two ways. First, the size of the commitment effect increases. Second, the impact of home price risk is reduced: in the limit where $\theta = 1$ and houses are never sold, home price risk becomes irrelevant (Sinai and Souleles, 2005). Higher adjustment costs unambiguously amplify the effect of housing on portfolio choice in environment where home price risk has little effect on portfolios (κ low). Hence, equation (2) captures the basic economic mechanisms connecting housing and stockholding that are emphasized in the literature.

Mortgage prepayment and borrowing constraints. The preceding analysis assumes that the size of the agent's mortgage is exogenous. In practice, mortgages can be pre-paid, raising the question of why households save in the form of liquid stocks and bonds while also holding mortgage debt. One natural reason may be that borrowing costs generate a motive to hold a buffer stock in liquid assets. We now show that the results on housing and portfolio choice obtained in Proposition 1 continue to hold in such an environment.

We generate buffer-stock savings behavior in a stylized manner by introducing a borrowing horizon $T_{BC} < T$. Agents are free to transfer labor income across periods up to T_{BC}, but cannot directly transfer income across periods before and after T_{BC}. In particular, the only

[5] Home price risk κ can also affect portfolios independently of this property value effect by creating fluctuations in the relative price of housing (Yogo 2006, Piazzesi et al 2007). This force alters the optimal portfolio in our model by changing the magnitude of the wealth effect (C_1).

way to borrow from periods after T_{BC} is through mortgage debt.[6] Let Y_{BC} denote labor income during periods up to T_{BC}.

We make the following assumptions about mortgage loans. At $t = 0$, the household is endowed with a home h_0, liquid wealth L_0 and pre-existing mortgage M_{-1}. The household can then choose to prepay or take on additional mortgage, resulting in an optimal mortgage choice $0 \leq M_0 \leq kp_0h_0$ where the upper bound reflects a limit on the size of mortgage relative to home value. A share θ of the loan is secured by commitment, and the rest by adjustable housing. Parelleling our assumption that commitment housing is unadjustable, we also assume that the mortgage over the commitment portion of the house cannot be altered after $t = 0$, and that a share T_{BC}/T of this committed mortgage is due during the first T_{BC} periods. In contrast, the mortgage over the adjustable part of housing is repaid at $t = 1$, and the consumer chooses adjustable housing for the next T_{BC} periods as a renter. The assumption that part of the mortgage is unadjustable is a stylized method of capturing the costs of borrowing against home equity.

Proposition 2 *With a borrowing horizon* $0 < T_{BC} < T$,

(i) The optimal mortgage M_0^* *at* $t = 0$ *is positive if* $(Y - Y_{BC}) / (Y + L_0 - M_{-1})$ *is large enough (holding fixed other model parameters).*

(ii) Given an initial mortgage M_0 *(whether chosen endogenously or exogenously), the optimal portfolio share of stocks out of liquid wealth at* $t = 0$ *is, to a log-linear approximation,*

$$\alpha^* = C_1 \cdot \frac{\textit{liquid wlth} + \textit{labor incm}_{BC} + [(1-\theta) + \theta T_{BC}/T] \cdot \textit{home eq}}{\textit{liquid wlth}} - \left[\theta C_1 \frac{T_{BC}}{T} + \kappa C_2\right] \cdot \frac{\textit{property val}}{\textit{liquid wlth}}$$

Part (i) of this proposition shows that it may be optimal to save in the form of liquid financial assets without fully paying off the mortgage. This is because the mortgage loan is used to overcome the constraint that borrowing against future labor income is not permitted. For a consumer who expects most of his income after T_{BC} (i.e., $(Y - Y_{BC}) / (Y + L_0 - M_{-1})$ is large) the mortgage loan is used for smoothing consumption over time. This consumer will simultaneously hold a mortgage and invest in liquid assets to finance consumption in the

[6] A weakness of this method of modelling borrowing constraints is that it prevents saving as well as borrowing: agents cannot transfer labor income earned before T_{BC} to finance consumption after T_{BC}. This restriction is less problematic when income is increasing over time.

first T_{BC} periods.[7]

Part (ii) shows that whether or not mortgage is set optimally, the portfolio allocation out of liquid wealth remains similar to that in the simple model. Agents with more expensive houses hold fewer risky assets in their buffer stock because of both the commitment (θ) and home price risk (κ) channels.

Estimating Equation. The key implication of the model is that one must distinguish changes in property value from changes in home equity wealth to fully uncover the effects of housing on portfolio choice. Linearizing the right hand side of equation (2) motivates an estimating equation for the stock share of liquid wealth of the form

$$\text{stock share}_i = \alpha + \beta_1 \text{property value}_i + \beta_2 \text{home equity}_i + \gamma X_i + \varepsilon_i \tag{3}$$

where X_i denotes a vector of controls such as liquid wealth, income, and other covariates. Under the null hypothesis that housing involves no commitment and no risk, $\beta_1 = 0$. Home equity and the other controls such as income or liquid wealth capture the wealth effect. The error term ε captures other sources of heterogeneity in portfolios. These may include entrepreneurial risk (Heaton and Lucas 2000), investment mistakes (Odean 1999, Calvet, Campbell and Sodini 2007), heterogeneity in risk aversion γ, or measurement error. Unfortunately, some of the effects captured by the error term may be correlated with property value, creating bias in an OLS estimate of β_1. One possible endogeneity problem, emphasized by Cocco (2005), stems from variation in unobserved labor income. To see how this affects the estimation, suppose that lifetime labor income is $Y = Y^{obs} + Y^{un}$, where Y^{obs} is observed by both the household and the econometrician, while Y^{un} is only observed by the household. The error term in (3) now becomes $\varepsilon = \beta_4 Y^{un} + \nu$, where ν represents other sources of heterogeneity. If property value is positively related to unobserved (e.g., future) labor income Y^{un} and labor income is positively associated with stock shares as in (2), we have $\mathrm{E}[\varepsilon \cdot \text{property value}] > 0$. In this case, the OLS estimate of β_1 is biased upward. Similarly, if households have heterogenous risk aversion γ and less risk averse households buy larger houses, we would again observe a spurious positive correlation between stockholdings

[7] In principle, the consumer could also finance expenditures by continuously increasing his mortgage debt and maintaining zero liquid wealth. In our stylized model, mortgage debt is fixed after $t = 0$, ruling out such behavior. More generally, costs of borrowing against home equity motivate consumers to maintain buffer stocks in the form of liquid wealth.

and property value. These problems make it essential to isolate variation in property value that is orthogonal to ε in order to identify β_1.

3 Data and Sample Definition

We estimate equation (3) using data from the 1990-2004 asset modules of the Survey of Income and Program Participation. Each SIPP panel tracks 20-30,000 households over a period of 2-3 years, collecting information on income, assets, and demographics. During the first four panels, asset data were only collected once; in the last three panels, asset data were collected once per year, permitting a panel analysis of changes in portfolios. The main advantages of the SIPP relative to other commonly used datasets on financial characteristics such as the SCF and PSID are its large sample size and detailed information about covariates such a complete housing history. We obtain quarterly data on average housing prices by state from 1975-2004 using the repeat sales index constructed by the Office of Federal Housing Enterprise Oversight (OFHEO).

The seven SIPP panels together contain information on 138,115 unique households. 70,857 of these households bought their current house after 1975 and therefore have OFHEO data for the year of home purchase, which is essential for our instrumental variable analysis. We exclude an additional 1,727 households whose total reported stockholdings are negative or exceed their liquid wealth. This exclusion does not affect the qualitative results reported below. These two exclusions leave us with 69,130 homeowners in our cross-sectional analysis sample. We observe asset data both before and after the purchase of a new house for 2,784 of these households. These households constitute our panel analysis sample.[8]

Table 1 reports summary statistics for the SIPP dataset as a whole, the cross-sectional analysis sample, and the panel analysis sample. In the cross-sectional sample, homeowners own houses that are worth approximately $122,000 on average in 1990 dollars. The average amount of home equity is $71,000 and the average outstanding mortgage is $51,000. The average household head is 47 years old and has lived in his current house for 8.5 years. Mean liquid wealth is $63,000, but this distribution is very skewed; the median level of liquid wealth is only $12,000. Mean total wealth (which includes liquid wealth, home equity, and other

[8] When we include these households in the cross-sectional sample, we only use data from the first year in which assets are observed. Hence, each observation in the cross-sectional sample is for a unique household.

illiquid assets such as cars) is $166,500. Households hold on average approximately 11% of their liquid wealth in the form of stocks in taxable (non-retirement) accounts and 40% in the form of "safe" assets (savings accounts, bonds, CDs, and money market accounts). The relatively small fraction of wealth held in stocks reflects the fact that only 27% of the households in the data hold stocks outside their retirement accounts (consistent with Vissing-Jorgensen 2002). Households hold nearly 30% of their liquid wealth in retirement accounts and 19% in other financial assets such as checking accounts, debt owed to the household, and equity in other financial investments. We do not have information on the risk characteristics of retirement portfolios or other financial assets, a limitation of the data that we address in our empirical analysis.

The summary statistics reported in Table 1 for the panel sample are for the year after home purchase. Homeowners in the panel sample generally have similar characteristics to those in the cross-sectional sample, with three exceptions. First, they have less home equity and more mortgage debt, as expected for new home buyers. Second, they are slightly less wealthy, consistent with being slightly younger on average. Finally, they hold more stocks in their portfolios. This is because the panel sample spans 1996-2003, a period with higher stock ownership than the early 1990s. These differences in sample characteristics should be kept in mind when comparing the empirical results obtained using the cross-sectional and panel identification strategies.

4 Empirical Analysis

To estimate (3), we seek variation in home equity and property value that is orthogonal to liquid wealth, wage income, and other unobserved determinants of portfolio choice (ε). We generate such variation using two approaches: a cross-sectional instrumental variables strategy and a panel strategy that permits household fixed effects in portfolio choice.

4.1 Cross-Sectional Results

Identification Strategy. The cross-sectional identification strategy exploits two instruments to generate variation in home equity and property value: the average price of houses in the individual's state in the current year (the year in which portfolios are measured) and the average price of houses in the individual's state in the year that he bought his house. The intuition for this identification strategy is illustrated in Figure 1, which plots average real home prices in California from 1975-2005 using the OFHEO data. Consider a hypothetical experiment involving a set of individuals who buy identical houses and only pay the interest on their mortgage (so that debt outstanding does not change over time). As a baseline, consider individual A who buys a house in 1985 (dashed red line) and whose portfolio we observe in 2000 (solid blue line), as shown in Panel A. Now compare this individual to individual B who buys the same house in 1990 and whose portfolio we also observe in 2000. Individuals A and B have the same current property value, but individual B is likely to have less home equity and a larger mortgage, because home prices were higher in 1990 than 1985. Intuitively, since individual B is buying the same house at a higher price, he needs a bigger mortgage; and because he enjoys less home price appreciation than A, he will end up with lower home equity in 2000. Now consider a second experiment, comparing panel C to A. Individual C buys the same house in 1985, but we observe his portfolio in 2005. This individual has the same mortgage debt as individual A (under the assumption that individuals only pay interest to service their debt), but has higher home equity and wealth at the time we observe his portfolio. Together, the two experiments (instruments) allow us to separately identify the causal effects of mortgages and home equity on portfolios.

In practice, our implementation of this strategy differs from the hypothetical examples in two ways. First, we do not just compare individuals who buy at different times, as such comparisons may be contaminated by time series fluctuations in asset prices or correlations between portfolios and home tenure or age. Because we have data on individuals who purchase houses in different years and observe portfolios in different years in 50 states, we include state, current year, year of house purchase, and age fixed effects in every regression specification below. Thus, we identify β_1 and β_2 in (3) purely from variation in house prices within state.[9] Second, unlike in the hypothetical example, individuals buy smaller houses when prices are high and reduce their mortgage debt over time by paying more than mortgage interest. The first stage effects of the house price indices on mortgage and home equity account for these effects.

[9] Campbell and Cocco (2007) show that financial market conditions are correlated with aggregate housing prices. Our empirical strategy is not biased by such correlations because we exploit differential variation across states in house prices within a single financial market.

Table 2 reports first stage regressions of mortgage, home equity, and property value (mortgage plus home equity) on the two instruments. The specifications in columns 1-3 include only state, year of purchase, current year, and age fixed effects as covariates. Column 1 shows that higher current house prices strongly predict higher property values, with a t-statistic of 42. Conditioning on current prices, higher house prices at the time of purchase predict slightly lower current property values, confirming that individuals purchase smaller houses if they buy at times when prices are relatively high. Column 2 shows that higher current prices strongly predict higher home equity, showing that much of the increase in property value comes from higher home equity, as expected. Higher prices at the time of purchase strongly predict lower home equity, with a t-statistic of 18. Conversely, column 3 shows that higher prices at the time of purchase predict much larger mortgages. Higher current prices also predict (to a smaller extent) larger mortgages, an effect that may be driven by refinancing – when current prices are high, individuals tap into their home equity. Finally, columns 4-6 show that these first-stage effects remain similar when we include a ten piece linear spline for liquid wealth as well as the following other controls: household income, household head's education, number of children, and the state unemployment rate in the current year as well as the year of home purchase. The results in Table 2 show that the two instruments meet the first condition for consistency of IV estimation, namely that they are powerful determinants of mortgage and home equity even conditional on a rich set of covariates.

The exclusion restriction for these instruments is that changes in average state house prices are orthogonal to unobserved determinants of portfolio decisions (ε). There are two potential threats to the validity of this exclusion restriction. First, fluctuations in local housing markets could be correlated with fluctuations in labor market or other economic conditions, which might in turn directly influence portfolio choices (an omitted variables problem). Second, the exclusion restriction could be violated via selection effects. People who buy houses when local prices are relatively high may have different risk preferences from those who buy when prices are lower. This could generate a spurious correlation between stock shares and house price indices. We address these threats to identification using a set of auxiliary tests as well as a separate panel identification strategy after presenting our baseline results.

Baseline Results. As a reference, we begin by estimating OLS specifications analogous to those in previous studies. Column 1 of Table 3 reports OLS estimates of a regression of the stock share of liquid wealth on home equity and property value without any covariates. The estimates imply that an increase in property value is strongly *positively* associated with the stock share of liquid wealth, contrary to the model's predictions. This is not surprising, as individuals with larger properties tend to be wealthier, and wealthier individuals hold more stocks. In Column 2 we control for wealth as well the full set of fixed effects and covariates used in Columns 4-6 of Table 2. The inclusion of these covariates reduces the coefficient on property value by an order of magnitude. These results echo the instability of OLS estimates found in prior studies. Moreover, they show that omitted variables and endogeneity problems tend to bias the effect of property value on stock shares upward, highlighting the need for improved identification to measure the negative causal effect of housing on stock shares.

The remaining columns of Table 3 report two-stage least squares estimates of β_1 and β_2 in (3), where home equity and property value are instrumented using the two OFHEO price indices. In column 3, we estimate the model including current year, year of purchase, age, and state fixed effects. The null hypothesis that changes in property value have no effect on financial portfolios is rejected with $p < 0.01$. The point estimate of the property value coefficient implies that a \$100,000 increase in an individual's mortgage reduces his stock share of liquid wealth by 6.59 percentage points. To interpret the magnitude of this coefficient, note that the mean stock share in the analysis sample is 11.3% and the mean mortgage is \$51,000. Hence, a 10% increase in mortgage debt (\$5,100) is estimated to reduce stock shares by 0.33 percentage points, and the elasticity of the stock share of liquid wealth with respect to mortgage debt is approximately -0.3. Equivalently, a \$10,000 increase in mortgage debt generates a 0.66 percentage point or roughly a 6% reduction in the portfolio share of stocks.

The estimate of the home equity coefficient in column 3 implies that a \$100,000 increase in home equity raises the stock share by 6.98 percentage points when total property value is held fixed, which we interpret as a wealth effect. The mean home equity in the sample is approximately \$71,000, implying an elasticity of stock share of liquid wealth with respect to home equity wealth of 0.44. Because property value is the sum of mortgage debt and home equity, our estimates for β_1 and β_2 imply that an increase in home equity holding fixed *mortgage debt* has no significant effect on portfolio allocations. This is because the wealth effect of having more home equity is cancelled out by the effect of owning a more expensive house. It is therefore crucial to disentangle the two components of property value in order to uncover the effects of housing on portfolios. An implication is that the demand for risky assets will not covary with current house price fluctuations (because they affect both wealth and property values simultaneously), but will covary negatively with outstanding mortgage debt.

Column 4 of Table 3 replicates column 3 with the full set of covariates in addition to the fixed effects – liquid wealth spline, education, income, number of children, and the state unemployment rate in the current year as well as the year of home purchase. The estimates of the coefficients of interest are virtually unaffected by controls, unlike the OLS estimates. Since controlling for observed heterogeneity has little impact on the estimates, one can be more confident that unobserved heterogeneity is unlikely to be driving the results. Moreover, because liquid wealth is held fixed in this specification, the results confirm that the changes in stock share are driven by the numerator (reductions in stocks) rather than changes in liquid wealth. Consistent with this result, we find using analogous regression specifications that increases in property value reduce dollars held in stocks but do not have a significant effect on liquid wealth (not reported).

In columns 5-6, we decompose the effects of housing on stock shares into stock market participation decisions and intensive margin changes in portfolio allocations. Column 5 replicates column 3, replacing the dependent variable with an indicator for owning stocks. A $100,000 increase in an individual's mortgage is estimated to reduce his probability of participating in the stock market by 14%, relative to a mean of 27%. Hence, the elasticity of stock market participation with respect to mortgage debt is approximately -0.25. Conversely, increases in home equity wealth increase the probability of stock market participation by a similar magnitude.

Column 6 isolates the intensive margin response – the change in stock shares conditional on participating in the stock market. This column reports estimates of a two-stage Tobit specification. This model is analogous to the two-stage-least-squares estimates, but corrects for the fact that some individuals are non-participants using a Tobit specification where the stock share is left censored at 0.[10] The estimates imply that a $100,000 increase in the

[10]Estimating a TSLS model only on the subsample of stock market participants yields biased estimates

mortgage outstanding reduces stock shares for stock market participants by 21% relative to a base of 39%. This implies an intensive-margin elasticity of stock shares with respect to mortgage debt of -0.25. Home equity changes again have similar effects in the opposite direction.

Robustness Checks. In Table 4, we evaluate the robustness of our estimates to alternative specifications and sample definitions. In column 1, we estimate a model analogous to Column 3 of Table 3 using logs instead of levels for the independent variables. We instrument for log(property value) and log(home equity) with the logs of the two OFHEO price indices. We retain the stock share in levels on the left hand side because of the large number of individuals with 0 stock shares in our sample. Consistent with the results in Table 3, the estimates reveal that increases in property value significantly reduce the share of stocks in liquid wealth while increases in home equity wealth increase stock shares.

Column 2 reports estimates from a specification where the endogenous regressors are also defined in shares, like the dependent variable. We replace property value by the ratio of property value to liquid wealth and home equity by the ratio of home equity to liquid wealth. We then use the level of the two OFHEO price indices as in Table 4 as instruments for these ratios. The advantage of this normalization is that we do not have to include liquid wealth as a separate regressor, which is preferable because it is an endogenous variable. The disadvantage is that we introduce substantial outliers: there are many observations with near-zero liquid wealth and large property values. To reduce noise from these outliers, we exclude observations with ratios of property value or home equity to liquid wealth above 30. The estimates are consistent with those obtained in Table 3, but less precisely estimated as expected because of the instability of the ratios.

In column 3, we replicate the baseline levels specification in column 3 of Table 3, but restrict the sample to individuals with more than $100,000 of total wealth. The objective of this specification is to assess whether the effects we have identified are also present among high-wealth households, whose behavior may be most relevant for financial market aggregates. The point estimate of the property value coefficient is similar in magnitude to that in the full sample, although it is less precisely estimated because of the reduced sample size ($p < 0.1$).

because changes in home equity and mortgages affect stock market participation rates, generating selection effects.

Housing remains an important determinant of portfolio choice even for wealthier households.

As noted above, one limitation of our dataset is that it does not contain information on the portfolio composition of retirement accounts. One may be concerned that the effects of housing on the portfolio composition of taxable accounts are offset by changes in portfolio allocations within retirement accounts. To alleviate this concern, in Column 4 we replicate the baseline specification for the subsample of households with zero retirement assets (approximately 45% of households). Reassuringly, the coefficients of interest are similar for this subsample – even households without retirement accounts shift their taxable asset holdings from stocks to safe assets when buying a more expensive house.

In Columns 5 and 6, we investigate where the money that is taken out of stocks goes. Column 5 replicates the baseline specification in Column 3 of Table 3, replacing the dependent variable with the share of wealth held in retirement accounts or "other financial assets." The estimates show that property values and home equity have little impact on the share of liquid wealth held in these assets. This is because the effects of housing on stock shares are fully offset by changes in the share of liquid wealth held in the form of taxable safe assets, as shown in Column 6. We conclude that individuals with larger houses shift their buffer stock from stocks to safe assets but do not significantly change the amount they allocate to retirement accounts or other financial assets.

Threats to Identification. Having established that there is a robust relationship between housing and portfolios, we now return to the two threats to identification discussed in the previous subsection. In order to evaluate these concerns, it is useful to understand the reduced-form relationships underlying the two-stage-least-squares estimates above. Two reduced-form relationships drive the results in Tables 3 and 4. First, individuals who buy houses when housing prices are relatively high in their state hold less stocks in subsequent years. Second, homeowners' stock shares do not vary substantially with contemporaneous housing prices. The first finding tells us that households with higher mortgage debt and lower home equity have lower stock shares. To determine which channel is responsible for the reduction in stockholding, we use the second finding, which shows that fluctuations in home equity have no effect on stock shares. This leads us to conclude that increases in mortgage debt reduce stockholding, as shown in Tables 3 and 4.

The first threat to causal interpretation of the two reduced-form relationships is that fluctuations in current home prices are correlated with local economic conditions that directly affect portfolio choice. We believe that such effects are unlikely to be responsible for our findings for three reasons. First, as shown above, controlling for the local business cycle using state unemployment rates has little effect on the estimates. Second, residual fluctuations in local house prices conditional on the set of fixed effects are not correlated with fluctuations in income or labor force participation. In particular, we find insignificant coefficients on property value and home equity in regressions that replace the dependent variable in specification 3 of Table 3 by income or an indicator for working. Third, most plausible omitted variable stories would bias the estimated effect of current house prices on stock shares upward. The finding that current house prices have much smaller effects on portfolios than house prices at the time of purchase is unlikely to be spuriously generated by such biases.

The second threat to identification is that fluctuations in house prices at the time of purchase are correlated with portfolios because of selection effects. Individuals who buy when house prices are relatively high may tend to hold safer portfolios simply because they have different risk preferences.[11] To evaluate whether this is the case, we turn to a panel analysis of homeowners for whom we observe portfolios both before and after the date of home purchase. If our results are driven by selection effects, individuals who buy when prices are high should hold more conservative portfolios even *before* they buy their houses.

4.2 Panel Results

Identification Strategy. The cross-sectional analysis in the preceding section controlled for unobserved differences across households in risk preferences using instruments for home equity and mortgage debt. In this section, we instead control for unobserved heterogeneity by allowing for household fixed effects using panel data. For an individual who buys a new house in year t, define $\Delta x = x_{t+1} - x_{t-1}$. Using the subsample of homeowners for whom we observe portfolio allocations before and after the purchase of a new house, we estimate (3)

[11] The most plausible form of selection is that individuals with lower risk aversion or less background risk buy houses when house prices are high. This would lead to a positive correlation between house prices in the year of purchase and stock shares. Such a correlation would bias the two-stage-least-squares estimates in Table 3 toward zero, again working against our findings.

in first differences:

$$\Delta\text{stock share}_i = \alpha + \beta_1 \Delta\text{property value}_i + \beta_2 \Delta\text{total wealth}_i + \gamma \Delta X_i + \Delta\varepsilon_i \qquad (4)$$

We first estimate (4) using OLS.[12] The identification assumption underlying this panel OLS analysis is that changes in property value are orthogonal to unobserved changes in households' risk preferences ($\Delta\varepsilon_i = \varepsilon_{i,t+1} - \varepsilon_{i,t-1}$).

The OLS panel analysis permits unobserved time-invariant differences across households in risk preferences, but may be biased by shocks to risk preferences that are correlated with the amount an individuals spends on a house. To address concerns about such biases, we estimate a second set of specifications where we instrument for Δproperty value using the state house price index in the year of home purchase. Because we only observe changes in portfolios over a short horizon, there is little difference between house prices at the time of purchase and the point at which we observe portfolio shares. Therefore, we cannot separately instrument for the effects of changes in wealth (via home equity) on portfolios as in the cross-sectional specifications. We find, however, that house price indices at the time of purchase are uncorrelated with changes in total wealth from the year before to the year after purchase. Thus, the IV estimate of β_1 in (4) is effectively identified from changes in property value that are orthogonal to changes in total wealth, as in the cross-sectional analysis.

OLS Results. Table 5 reports OLS estimates of several variants of (4). To reduce the influence of outliers, we exclude 33 households (1% of the sample) who report changes in total wealth of more than \$1 million in magnitude from the year before to the year after home purchase. Column 1 of Table 5 shows estimates of (4) without any additional controls. The estimates imply that a \$100,000 increase in property value leads to a 3.6 percentage point reduction in the stock share of liquid wealth, an estimate that is statistically significant with $p < 0.01$. In Column 2, we include the following set of additional controls: state, year of purchase, and age fixed effects, education, number of children, and state unemployment rate in the year before home purchase, and change in household income from the year before to the year after home purchase. The inclusion of these covariates does not significantly affect

[12] We control for total wealth in (4) rather than separating home equity from liquid wealth because there is insufficient power to separately identify the effects of changes in property value and home equity in the panel analysis. The reason is that changes in home equity and property value are very highly correlated from the year before to the year after home purchase, the horizon over which we observe changes in portfolios.

the estimated effect of property value on the stock share. This result confirms our prior that panel OLS estimates identified from within-household changes suffer less from omitted variable bias than cross-sectional OLS regressions that compare different households.

In columns 3-5 of Table 5, we explore the mechanism through which households reduce stock shares of liquid wealth when buying a house. In column 3, we replicate the specification in column 1 using the change in liquid wealth as the dependent variable. The coefficient estimate implies that buying a $100,000 more expensive house leaves a household with $33,210 less in liquid wealth, presumably because of the need to finance a large downpayment. Columns 4 and 5 investigate the composition of this reduction in liquid wealth using the change in dollars held in stocks and safe assets as dependent variables. Column 4 shows that households finance most of their downpayments by selling stocks – a $100,000 increase in property value leads to a $26,505 reduction in assets held in stocks. In contrast, the same increase in property value leads to only a $2,356 reduction in safe assets. The finding that individuals tend to sell stocks rather than bonds in order to finance the downpayment for a more expensive house constitutes direct evidence that housing induces individuals to take less risk in their buffer stock of liquid wealth.

IV Results. Next, we estimate (4) using the state house price as an instrument for property values. Because the panel sample spans a short period – 1996 to 2003 – and has much fewer observations, there is inadequate power to include both state and year of purchase fixed effects in this analysis. Therefore, we include only state and age fixed effects in most specifications. We then show that similar but less precisely estimated results are obtained when we also permit year effects.

Columns 1-3 of Table 6 document the first-stage effects of the state house price index on changes in property value, home equity, and mortgage debt in a regression that includes state and age fixed effects as well the change in total wealth. As expected, individuals who buy houses in higher priced markets spend more on their houses, with significant increases in both home equity and mortgage debt. The price index remains a highly significant predictor of changes in property values when we include the full set of controls used in Column 2 of Table 5, with the exception of year fixed effects (not reported).

Columns 4-6 report TSLS estimates of the effect of changes in property value on the stock share of liquid wealth. In column 4, we include only state and age fixed effects and the change in total wealth as controls. A $100,000 increase in property value is estimated to reduce the stock share by 14.6 percentage points in this specification. This estimate is statistically significant with $p < 0.05$. Column 5 shows that controlling for education, number of children, state unemployment rate, and the change in household income does not affect this estimate significantly. Finally, column 6 also introduces year fixed effects. As expected, we lose precision because there is little differential variation in house prices across states over the few years in this sample. Nevertheless, it is reassuring that the point estimate of the coefficient remains relatively stable and negative.

It is not surprising that the IV panel estimates are more negative than the OLS panel estimates. Even the panel OLS estimates are likely to be biased upward because of the endogeneity of housing choice. For instance, if individuals move to more expensive houses when their background risk falls, the OLS estimate of β_1 would be biased even in the panel. The IV panel estimates are also roughly twice as large as the IV cross-section estimates reported in Tables 3 and 4. This is most likely because the mean stock share in the panel analysis sample is significantly larger than that in the cross-sectional sample, as shown in Table 1. The panel IV strategy implies that a $10,000 increase in mortgage debt would reduce the stock share by 9%, only slightly larger than the corresponding cross-sectional IV estimate of 6%.

The panel analysis shows that the difference in portfolios between individuals who buy when house prices are high and low emerges immediately *after* home purchase, suggesting that these differences are caused by the house itself and not intrinsic variation in risk preferences across individuals. This finding helps mitigate concerns that selection and omitted variable issues are responsible for our findings in the cross-sectional IV analysis. Together, the cross-sectional and panel identification strategies show that housing purchases have both immediate and long-lasting effects on household's portfolio choices.

4.3 Commitments vs. House Price Risk

As noted above, the theoretical literature has identified two reasons why housing may induce households to hold safer portfolios: (1) greater exposure to house price risk and (2) a larger consumption commitment. We now present some suggestive evidence to distinguish between these two channels. The tests we implement are motivated by the comparative statics of the portfolio choice equation in (2). We examine how the effect of housing on portfolios covaries with the volatility of the local housing market and the household's adjustment costs. The house price risk channel predicts that owning a larger house should have a larger effect on stock shares in riskier markets, while the commitment channel predicts that the effect should be larger for households that face higher adjustment costs.

Columns 1 and 2 of Table 7 test for differential effects of housing on portfolio choice by the riskiness of the local housing market. To maximize power, we use the cross-sectional analysis sample in this section. We estimate specifications that include the full set of controls used above to control for potential differences in household characteristics across low-risk and high-risk housing markets. We first compute the standard deviation of annual house price growth rates using the OFHEO data by state. States with above-median volatility (standard deviation $> 4.5\%$) are classified as "high risk," while the remainder are classified as "low risk." To test whether the effects of housing on portfolios differ in high vs. low risk environments, we interact the high risk indicator with property value and home equity and estimate a specification analogous to (3). We instrument for the interaction effects using the interactions of the two OFHEO price indices used in section 4.1 and the high risk indicator. We use the same set of covariates as in Column 4 of Table 3 and also include the high risk indicator directly. The coefficient estimates in column 1 imply that a $100,000 increase in property value reduces the stock share of liquid wealth by 9 percentage points. The coefficient on the interaction between property value and the high risk indicator has a coefficient of $+1.5\%$ and is statistically insignificant. Similarly, the interaction of home equity and the risk indicator is near zero and insignificant. Hence, housing has a similar effect on portfolios in low and high risk states. Column 2 replicates the specification in column 1 but replaces the dependent variable with an indicator for owning stocks. Consistent with the stock share results, housing has a substantial and statistically significant effect on stock market participation, but does not have a differential effect in low- and high-risk states.

Next, we test for differential effects of housing on portfolios by the household's adjustment costs. Because we do not directly observe adjustment costs, we proxy for them using mean home tenures. Following the same strategy as for the home price risk test above, we first compute the mean home tenure in each state in our data. States with mean home tenures above the sample median (mean home tenure > 8.4 years) are classified as "high adjustment cost." Columns 3 and 4 replicate columns 1 and 2 of Table 7, replacing the high risk indicator with the high adjustment cost indicator. The interaction effects are instrumented for using interactions of the two OFHEO price indices and the high adjustment cost indicator. A $100,000 increase in property value reduces the stock share by 5.9 percentage points for households living in low adjustment cost states vs. 8.4 percentage points for those living in high adjustment cost states. Similarly, changes in home equity also have larger effects on portfolios for households with high adjustment costs. The interactions of home equity and property value with the high adjustment cost indicator are both significant with $p < 0.05$. Column 4 shows that results for stock market participation are similar: owning a more expensive house reduces the probability of stock market participation significantly for more households in states with higher adjustment costs.

Although not conclusive because the variation in house price risk or adjustment costs is not exogenous, these auxiliary tests suggest that housing induces households to hold safer portfolios primarily because of the commitment effect.[13] House price risk may have small effects on portfolios because most households are well hedged against house price risk (Sinai and Souleles 2005). This evidence suggests that housing may induce individuals to invest less in stocks particularly during periods when the housing market is illiquid and mobility rates are low.

5 Conclusion

This paper has characterized the causal effect of housing on portfolio choice. We find that an increase in mortgage debt, holding wealth fixed, reduces a household's propensity to participate in the stock market and reduces the share of stocks in the portfolio conditional on participation. The estimated elasticity of the share of liquid wealth allocated to stocks with respect to mortgage debt is -0.3. Increases in home equity wealth while holding property value fixed increase stockholding. The estimated elasticity of the stock share of liquid wealth with respect to home equity is 0.44. These elasticities are larger for households with larger adjustment costs, but similar across high and low-risk housing markets. Overall, the evidence confirms the prediction of theoretical models and simulations such as that of Flavin

[13] The model predicts that higher adjustment costs will amplify the coefficient on housing precisely when the effect of home price risk on portfolios (κ) is small, and is thus consistent with the results of the two tests.

and Yamashita (2002) and Cocco (2005), which predict that housing has substantial effects on financial portfolios. Our empirical results show that housing should indeed be a central element of portfolio choice and asset price models.

The evidence in this paper suggests that the interaction between housing and financial markets could have important consequences for the macroeconomy. In the recent past, there have been three rapid changes in housing markets: a substantial increase in mortgage debt, a rapid decline in property values, and a substantial increase in the illiquidity of housing as many individuals postpone selling their homes. Our empirical evidence suggests that each of these factors induces households to withdraw funds from stocks. Hence, recent changes in the housing market could potentially have reduced the demand for risky assets and exacerbated the decline in financial markets. In future work, it would be interesting to explore whether such interactions are consistent with historical fluctuations in housing and asset prices using calibrated general equilibrium models. Our analysis also suggests that the illiquidity of housing amplifies the welfare costs of shocks, making households more risk averse. Reducing transaction costs in the housing market could raise individual welfare both directly and by allowing households to choose better financial portfolios.

Appendix

Proof of Proposition 1. The total amount of resources available at $t = 1$ is

$$W_1 = L_1 + Y + p_1 h_0 - M$$

where L_1 is liquid wealth, Y is labor income, p_1 is the price per unit of housing, and M is mortgage. Out of this wealth, the consumer has committed buying Tx_0 units of the commitment good, which has a dollar cost of $T \cdot (p_1/T)\, h_0\theta$ since p_1/T is the per period price of housing. Thus, disposable wealth that can be spent on adjustable consumption is

$$W_1^D = L_1 + Y + p_1 h_0 - M - T \cdot (p_1/T)\, h_0\theta.$$

The present value of this risky payoff at date zero is

$$W_0^D = L_0 + Y + p_0 h_0 - M - p_0 h_0 \theta$$

where L_0 is initial liquid wealth. The consumer's investment problem can now be broken into two parts. First, to finance commitment consumption, the consumer holds a portfolio that generates $p_1 h_0 \theta$ in date $t = 1$. This requires dollar stockholdings of $\kappa\theta p_0 h_0$ because of our assumption that $p_1/p_0 = 1 + \kappa R$.

Second, the consumer invests optimally out of W_0^D to finance adjustable consumption. We now solve for this optimal choice. Denote the relative price of renting by $q = p_1/T$. Cobb-Douglas preferences and the absence of discounting imply that the optimal choice of f and y in every period are

$$f_t = \frac{W_1^D}{T} \cdot \frac{\mu_f}{\mu_f + \mu_a} \text{ and } y_t = \frac{1}{q}\frac{W_1^D}{T} \cdot \frac{\mu_a}{\mu_f + \mu_a}.$$

Noting that $q = (1 + \kappa R)\, p_0/T$, the value function over adjustable consumption is proportional to

$$\frac{\left(W_1^D\right)^{\mu_f + \mu_x}}{(1 + \kappa R)^{\mu_x}}$$

where the denominator represents the relative price risk.

Now write $W_1 = (1 + \alpha R)\, W_0$ where W_0^D to find the optimal portfolio share out of disposable wealth α. We need to maximize

$$\frac{(1+\alpha R)^{\left(\mu_f+\mu_x\right)(1-\gamma)}}{(1+\kappa R)^{\mu_x(1-\gamma)}}.$$

Following Campbell and Viceira (2002), we solve this optimization problem using lognormal approximations. Assuming lognormality of both $1+\alpha R$ and $1+\kappa R$, our maximization problem is equivalent to

$$\max_{\alpha} \mathrm{E} r_p + \frac{1}{2}(1-\gamma)\left(\mu_f+\mu_x\right)\sigma_p^2 - (1-\gamma)\,\mu_x \mathrm{cov}\left[r_p, r_h\right]$$

where $r_h = \log(1+\kappa R)$. Using the approximations

$$\begin{aligned} r_p &= \alpha r + \frac{1}{2}\alpha(1-\alpha)\sigma^2 \\ r_h &= \kappa r + \frac{1}{2}\kappa(1-\kappa)\sigma^2 \end{aligned}$$

of Campbell and Viceira, we obtain a linear-quadratic problem in α which yields

$$\alpha = \frac{\left[\mathrm{E}r + \frac{1}{2}\sigma^2\right] - (1-\gamma)\,\mu_x \cdot \kappa\sigma^2}{\sigma^2\left[1-(1-\gamma)\left(\mu_f+\mu_x\right)\right]}.$$

The optimal dollar amount invested in stocks is the sum of stocks used to finance adjustable and housing consumption, i.e., $\alpha W_0^D + \kappa\theta p_0 h_0$. However, the household is already holding some stocks implicitly in the form of the house. This stock exposure corresponds to a implicit dollar stockholding of $\kappa p_0 h_0$. Thus the optimal portfolio share of stocks out of liquid wealth is

$$\alpha^* = \alpha\frac{L_0 + Y + p_0 h_0 - M}{L_0} - \alpha\theta\frac{p_0 h_0}{L_0} - \kappa(1-\theta)\frac{p_0 h_0}{L_0}$$

which is the expression given in the proposition, with $C_1 = \alpha$ and $C_2 = 1-\theta$.

Proof of Proposition 2. We start with (ii), assuming that M_0 is exogenously fixed. The total amount of disposable resources available at $t=1$, net of commitment consumption, is

$$W_1^D = L_1 + Y_{BC} + (1-\theta)\left[p_1 h_0 - M_0\right] - \theta M_{BC}$$

where Y_{BC} is labor income up to T_{BC}, and $M_{BC} = (T_{BC}/T)\cdot M_0$ is mortgage up to T_{BC}.

The present value of this at $t = 0$ is

$$W_0^D = L_0 + Y_{BC} + (1-\theta)[p_0 h_0 - M_0] - \theta M_{BC}.$$

Optimal stockholdings are a share α of this, but the agent already holds $\kappa(1-\theta)p_0 h_0$ in the form of adjustable housing, resulting in

$$\begin{aligned} \alpha^* &= \alpha \frac{L_0 + Y_{BC} + (1-\theta)[p_0 h_0 - M_0] - \theta (T_{BC}/T) \cdot M_0}{L_0} - \kappa(1-\theta)\frac{p_0 h_0}{L_0} \\ &= \alpha \frac{L_0 + Y_{BC} + (1-\theta+\theta(T_{BC}/T))[p_0 h_0 - M_0]}{L_0} - [\kappa(1-\theta) + C_1 \theta (T_{BC}/T)]\frac{p_0 h_0}{L_0} \end{aligned}$$

as desired.

Now we show (i). Lifetime utility is the sum of expected utility in the first T_{BC} periods, plus expected utility in the last $T - T_{BC}$ periods. Given the level of adjustable housing $(1-\theta)p_0 h_0$, maximized expected utility in the first T_{BC} periods depends only on available cash resources

$$L_0 + Y_{BC} - M_{-1} + \theta[1 - T_{BC}/T] \cdot M_0^*.$$

To understand this, note that if mortgage is fully repaid, available cash wealth is $L_0 + Y_{BC} - M_{-1}$. On top of that, each dollar of mortgage outstanding increases available cash by $\theta[1 - T_{BC}/T]$ dollars, because this is the share of the mortgage that is on commitment housing and that will be due after T_{BC}. Similarly, maximized expected utility after T_{BC} depends only on available cash $Y - Y_{BC} - \theta(1 - T_{BC}/T) M_0^*$. We can now write total expected utility as

$$U^1(L_0 - M_{-1} + Y_{BC} + \theta[1 - T_{BC}/T] \cdot M_0^*) + U^2(Y - Y_{BC} - \theta(1 - T_{BC}/T) M_0^*)$$

where both U^1 and U^2 are concave indirect utility functions. It follows that maximizing lifetime expected utility in M_0^* is a concave problem with a unique optimum in the interval $[0, kp_0 h_0]$. Note that $(Y - Y_{BC}) / (Y + L_0 - M_{-1})$ is the ratio of resources spent on adjustable consumption after T_{BC} relative to total income assuming that all mortgage is prepaid. Consumption smoothing dictates that when this ratio is large, the consumer wants to transfer resources from the future to the present, resulting in a positive M_0^*.

References

Brueckner, J.K. (1997). "Consumption and Investment Motives and the Portfolio Choices of Homeowners," *Journal of Real Estate Finance and Economics*, 15(2): 159–80.

Calvet, L.E., Campbell, J.Y. and Sodini, P. (2007). "Down or Out: Assessing the Welfare Costs of Household Investment Mistakes," *Journal of Political Economy*, 115(5): 707–747.

Campbell, J. Y. and Cocco, J. F. (2003). "Household Risk Management and Optimal Mortgage Choice," *Quarterly Journal of Economics* 118:1449-1494.

Campbell, J. Y. and Cocco, J. F. (2007). "How Do House Prices Affect Consumption? Evidence from Micro Data," *Journal of Monetary Economics* 54:591-621.

Campbell, J. Y. and Viceira, L. M. (2002) *Strategic Asset Allocation: Portfolio Choice for Long Horizon Investors*, Oxford University Press.

Chetty, R. and Szeidl, A. (2007). "Consumption Commitments and Risk Preferences," *Quarterly Journal of Economics* 122(2): 831-877.

Cocco, J. (2005). "Portfolio Choice in the Presence of Housing," *Review of Financial Studies* 18, 535–567.

Davidoff, Thomas (2009). "Home Equity Commitment and Long-Term Care Insurance Demand," forthcoming *Journal of Public Economics.*

Flavin, M. and Yamashita, T. (2002). "Owner-Occupied Housing and the Composition of the Household Portfolio," *American Economic Review*, 92(1): 345–362.

Fratantoni, M.C. (2001). "Homeownership, Committed Expenditure Risk, and the Stockholding Puzzle," *Oxford Economic Papers*, 53(2): 241–59.

Fratantoni, M.C. (1998). "Homeownership and Investment in Risky Assets," *Journal of Urban Economics*, 44(1): 27–42.

Grossman, S. J. and Laroque, G. (1990). "Asset Pricing and Optimal Portfolio Choice in the Presence of Illiquid Durable Consumption Goods," *Econometrica* 58, 25–51.

Heaton J. and Lucas, D.J. (2000). "Portfolio Choice in the Presence of Background Risk," *Economic Journal* 110(460): 1-26.

Odean, T. (1999). "Do Investors Trade Too Much?" *American Economic Review*, 89(5): 1279–1298

Piazzesi, M., Schneider, M. and Tuzel, S. (2007). "Housing, Consumption and Asset Pricing," *Journal of Financial Economics*, 83(3): 531–569.

Shore, S. and Sinai, T. (2009). "Commitment, Risk, and Consumption: Do Birds of a Feather Have Bigger Nests?" *Review of Economics and Statistics*, forthcoming.

Sinai, T. and Souleles N.S. (2005). "Owner-Occupied Housing as a Hedge Against Rent Risk," *Quarterly Journal of Economics,* 120(2): 763–789.

Vissing-Jorgensen, A. (2002), "Towards an Explanation of Household Portfolio Choice Heterogeneity: Nonfinancial Income and Participation Cost Structures," NBER Working Paper 8884.

Yamashita, T. (2003). "Owner-Occupied Housing and Investment in Stocks: An Empirical Test," *Journal of Urban Economics*, 53(2): 220–237.

Yogo, M. (2006). "A Consumption-based Explanation of Expected Stock Returns," *Journal of Finance*, 61(2): 539–58.

TABLE 1
Summary Statistics for SIPP Estimation Samples

	Full Sample (1)	Cross-Sectional Analysis Sample (2)	Panel Analysis Sample (3)
Demographics:			
Age (years)	49.09	47.74	44.14
Years of education	12.69	13.52	13.58
Number of children	0.55	0.66	0.95
Household Income ($)	34,275	46,890	49,413
Housing:			
Property value ($)	82,484	122,678	123,286
Home equity ($)	45,090	71,220	54,348
Mortgage ($)	25,681	51,457	69,499
Home tenure (years)	14.93	8.49	1.00
Wealth:			
Liquid wealth ($)	41,328	62,634	57,220
Median liquid wealth ($)	3,581	12,224	6,583
Total wealth ($)	108,394	166,533	139,549
Median total wealth ($)	42,990	89,932	60,396
Portfolio Allocation:			
Percent of households holding stock	19.17%	27.23%	31.36%
Stock share (% of liquid wealth)	9.51%	11.32%	16.43%
Safe assets share (% of liquid wealth)	44.77%	40.44%	41.31%
Retirement assets share (% of liquid wealth)	22.59%	29.46%	23.19%
Other fin. assets share (% of liquid wealth)	23.13%	18.77%	19.07%
Number of observations	138,115	69,130	2,784

Notes: Table entries are means unless otherwise noted. Column 1 is based on all household heads (reference persons) in the 1990-2006 SIPP panels. Column 2 includes all household heads who purchased houses in or after 1975 and for whom house price index information is available, which is the estimation sample for the cross-sectional analysis. Column 3 includes only the subset of household heads in the 1996 and 2001 SIPP panels for whom we observe wealth both before and after the year of home purchase, which is the estimation sample for the panel analysis. All monetary values are in real 1990 dollars. Home tenure is defined as numbers of years living in current house.

TABLE 2
House Price Indices: First Stage Regression Estimates

Dependent Variable.:	Prop Val ($) (1)	Home Equity ($) (2)	Mortgage ($) (3)	Prop Val ($) (4)	Home Equity ($) (5)	Mortgage ($) (6)
OFHEO state house price index in current year	377.62	326.04	51.58	345.79	296.34	49.45
	(9.06)	(7.61)	(4.97)	(8.70)	(7.73)	(5.07)
	[41.69]	[42.85]	[10.37]	[39.740]	[38.35]	[9.76]
OFHEO state house price index in year of purchase	-54.52	-186.79	132.27	-42.81	-171.79	128.98
	(11.77)	(9.88)	(6.46)	(10.93)	(9.70)	(6.36)
	[4.63]	[18.9]	[20.47]	[3.92]	[17.71]	[20.27]
state, current year, purch. year and age FE's	x	x	x	x	x	x
liquid wealth spline				x	x	x
other controls				x	x	x
Observations	69,130	69,130	69,130	67,034	67,034	67,034

Notes: Standard errors in parentheses and t-statistics in square brackets. Specifications 1-3 include fixed effects for the household head's state of residence, age, current year (year in which portfolio allocations and current property value are measured), and year of home purchase. Specifications 4-6 include these fixed effects, a 10-piece linear spline for liquid wealth, and the following other controls: income, education, number of children, and state unemployment rate in the current year as well as the year of home purchase. In columns 1 and 4, the dependent variable is property value in the current year; in 2 and 5, it is home equity in the current year; and in 3 and 6, it is total outstanding mortgage debt in the current year. All monetary values are in real 1990 dollars.

TABLE 3
Cross-Sectional OLS and Instrumental Variable Estimates

	OLS	OLS	Two-Stage Least Squares			Two-step Tobit
Dependent Variable:	Stock Share (%)	Stock Share (%)	Stock Share (%)	Stock Share (%)	Stock Holder (%)	Stock Share (%)
	(1)	(2)	(3)	(4)	(5)	(6)
Property val.	4.87	0.58	-6.59	-6.90	-13.87	-20.84
(x $100K)	(0.21)	(0.21)	(2.19)	(2.14)	(4.12)	(7.68)
Home equity	-0.89	-0.62	6.98	4.84	14.69	24.45
(x $100K)	(0.25)	(0.24)	(2.53)	(2.52)	(4.78)	(8.77)
state, curr. year, purch. year and age FE's	x	x	x	x	x	x
liquid wealth spline		x		x		
other controls		x		x		
Observations	69,130	67,034	69,130	67,034	69,130	69,130

Notes: Standard errors in parentheses. All specifications include fixed effects for the household head's state of residence, age, current year (year in which portfolio allocations and current property value are measured), and year of home purchase. Specification 2 and 4 include these fixed effects, a 10-piece linear spline for liquid wealth, and the following other controls: income, education, number of children, and state unemployment rate in current year as well as year of home purchase. The dependent variable in specifications 1-4 and 6 is dollars held in stocks divided by liquid (non-housing) wealth. The dependent variable in specification 5 is an indicator for holding stocks. Specifications 1-2 are estimated using OLS; 3-5 are estimated using two-stage least squares; and 6 is estimated as a Tobit model with endogenous regressors using Newey's two-step estimator. Instruments for property value and home equity are the current-year and year of purchase OFHEO state price indices in all specifications. Coefficients can be interpreted as percentage point effect of a $100,000 change in property value and home equity.

TABLE 4
Cross-Sectional IV Estimates: Robustness Checks

Dependent Variable:	Stock Share of Liquid Wealth				Other Asset	Safe Asset
	Logs	Shares	High Wealth	No IRA/401k	Share	Share
	(%)	(%)	(%)	(%)	(%)	(%)
	(1)	(2)	(3)	(4)	(5)	(6)
log prop value	-26.40					
(x $100K)	(10.46)					
log home equity	12.24					
(x $100K)	(5.34)					
Prop val/liq wealth		-4.50				
(x $100K)		(2.20)				
Home eq/liq wlth		3.57				
(x $100K)		(1.74)				
Property value			-6.52	-8.20	0.47	6.79
(x $100K)			(4.01)	(3.52)	(3.50)	(3.44)
Home equity			9.61	9.42	2.28	-8.35
(x $100K)			(5.01)	(3.95)	(4.05)	(3.99)
state, curr. year, purch. year and age FE's	x	x	x	x	x	x
Observations	69,130	47,680	32,057	36,556	69,130	69,130

Notes: Standard errors in parentheses. All specifications report two-stage-least square estimates and include fixed effects for state of residence, age, current year, and year of home purchase. Instruments for property value and home equity are the current-year and year of purchase OFHEO state price indices in specifications 2-6; in specification 1, we instrument for the logs of these variables with the logs of the price indices. In specification 2, the endogenous regressors are the ratio of property value to liquid wealth and the ratio of home equity to liquid wealth; households for whom either of these ratios exceed 30 are excluded in this specification. Specification 3 replicates the baseline specification in column 3 of Table 3, restricting the sample to individuals with total wealth above $100,000. Specification 4 instead restricts the sample to those who report zero wealth in retirement accounts. Specifications 5 and 6 replicate the baseline specification on the full sample with different dependent variables. In specification 5, the dependent variable is total assets held in retirement accounts and other financial assets divided by liquid wealth. In specification 6, the dependent variable is safe assets held in taxable accounts divided by liquid wealth.

TABLE 5
OLS Panel Estimates

Dependent Variable:	Δ stock share (%) (1)	Δ stock share (%) (2)	Δ $ liq. wealth ($) (3)	Δ $ stocks ($) (4)	Δ $ safe assets ($) (5)
Δ Property value	-3.61	-3.51	-33,210	-26,505	-2,356
(x $100K)	(0.70)	(0.75)	(1483)	(1438)	(798)
Δ total wealth	x	x	x	x	x
state, year, age FE's		x	x	x	x
Other controls		x			
Observations	2,243	2,211	2,750	2,750	2,750

Notes: Standard errors in parentheses. All specifications report OLS estimates of the effect of changes in property value from the year before to the year after home purchase using the panel analysis sample. Property value is defined as 0 in the year before home purchase for those who were previously renting. All specifications control for the change in total household wealth from the year before to the year after home purchase. Specifications 2-5 include fixed effects for state of residence, age, and year of home purchase. Specification 2 also includes the following other controls: change in income from year before to year after home purchase as well as education, number of children, and state unemployment rate in year before home purchase. The dependent variable in specifications 1 and 2 is change in the stock share of liquid wealth from the year before to the year after home purchase. The dependent variable is change in liquid wealth from year before to year after home purchase in specification 3, changes in dollars held in stocks in column 4, and change in dollars held in safe assets in column 5.

TABLE 6
Instrumental Variable Panel Estimates

	First Stage (OLS)			Two-Stage Least Squares		
Dependent Variable:	Δ Prop Val (%)	Δ Home Eq (%)	Δ Mortg (%)	Δ Stock Share of Liquid Wealth (%)	(%)	(%)
	(1)	(2)	(3)	(4)	(5)	(6)
OFHEO state house price index in year of purchase	276.52 (49.19)	119.43 (38.83)	152.29 (30.80)			
Δ Property value (x $100K)				-14.63 (6.97)	-14.75 (6.56)	-19.29 (14.13)
state, age FE's	x	x	x	x	x	x
Δ total wealth	x	x	x	x	x	x
other controls					x	x
year FE's						x
Observations	2,753	2,753	2,753	2,248	2,243	2,243

Notes: Standard errors in parentheses. Specifications 1-3 report OLS estimates of the first-stage effect of the house price index in the year of purchase on the change in property value, home equity, and mortgage debt from the year before to the year after home purchase. Specifications 4-6 report TSLS estimates of the effect of changes in property value on changes in the stock share of liquid wealth using this instrument. All specifications include fixed effects for state of residence and age and also control for the change in total wealth from the year before to the year after home purchase. Specifications 5 and 6 also include the following other controls: change in income from year before to year after home purchase as well as education, number of children, and the state unemployment rate in year before home purchase. Specification 6 also includes fixed effects for the year of home purchase.

TABLE 7
House Price Risk vs. Commitment Effects

	Price Risk Interactions		Adj. Cost Interactions	
Dependent Variable:	Stock Share (%) (1)	Stockholder (%) (2)	Stock Share (%) (3)	Stockholder (%) (4)
Property value (x $100K)	-8.99	-19.02	-5.88	-12.39
	(2.59)	(4.54)	(2.02)	(3.54)
Home equity (x $100K)	5.88	12.65	3.66	8.05
	(2.94)	(5.15)	(2.45)	(4.28)
High risk x prop val (x $100K)	1.53	3.05		
	(1.11)	(1.94)		
High risk x home equity (x $100K)	-0.47	-0.83		
	(1.22)	(2.13)		
High adj cost x prop value (x $100K)			-2.47	-6.14
			(1.17)	(2.05)
High adj cost x home equity (x $100K)			2.57	5.81
			(1.17)	(2.06)
Full controls	x	x	x	x
High risk indicator	x	x		
High adjustment cost indicator			x	x
Observations	67,034	67,034	67,034	67,034

Notes: Standard errors in parentheses. All specifications report TSLS estimates where the instruments for property value and home equity are the current-year and year of purchase OFHEO state price indices. All specifications include the following set of full controls: fixed effects for state of residence, age, current year, and the year of home purchase, as well as a 10-piece linear spline for liquid wealth, income, education, number of children, and the state unemployment rate in the current year as well as the year of home purchase. In specifications 1 and 2, high risk is an indicator for whether the individual lives in a state whose standard deviation of house prices as computed based on the OFHEO index exceeds the median. These specifications include the high risk indicator and interactions of high risk with the property value and home equity variables. The instruments for the interactions are the interactions of the two price indices with the high risk indicator. Specifications 3 and 4 replicate 1 and 2, replacing the high risk indicator with a high adjustment cost indicator. High adjustment cost is an indicator for whether the individual lives in a state whose mean home tenure exceeds the median value in the sample. The dependent variable is the stock share of liquid wealth in specifications 1 and 3 and an indicator for holding stocks in specifications 2 and 4.

FIGURE 1
Real Housing Prices in California, 1975-2005

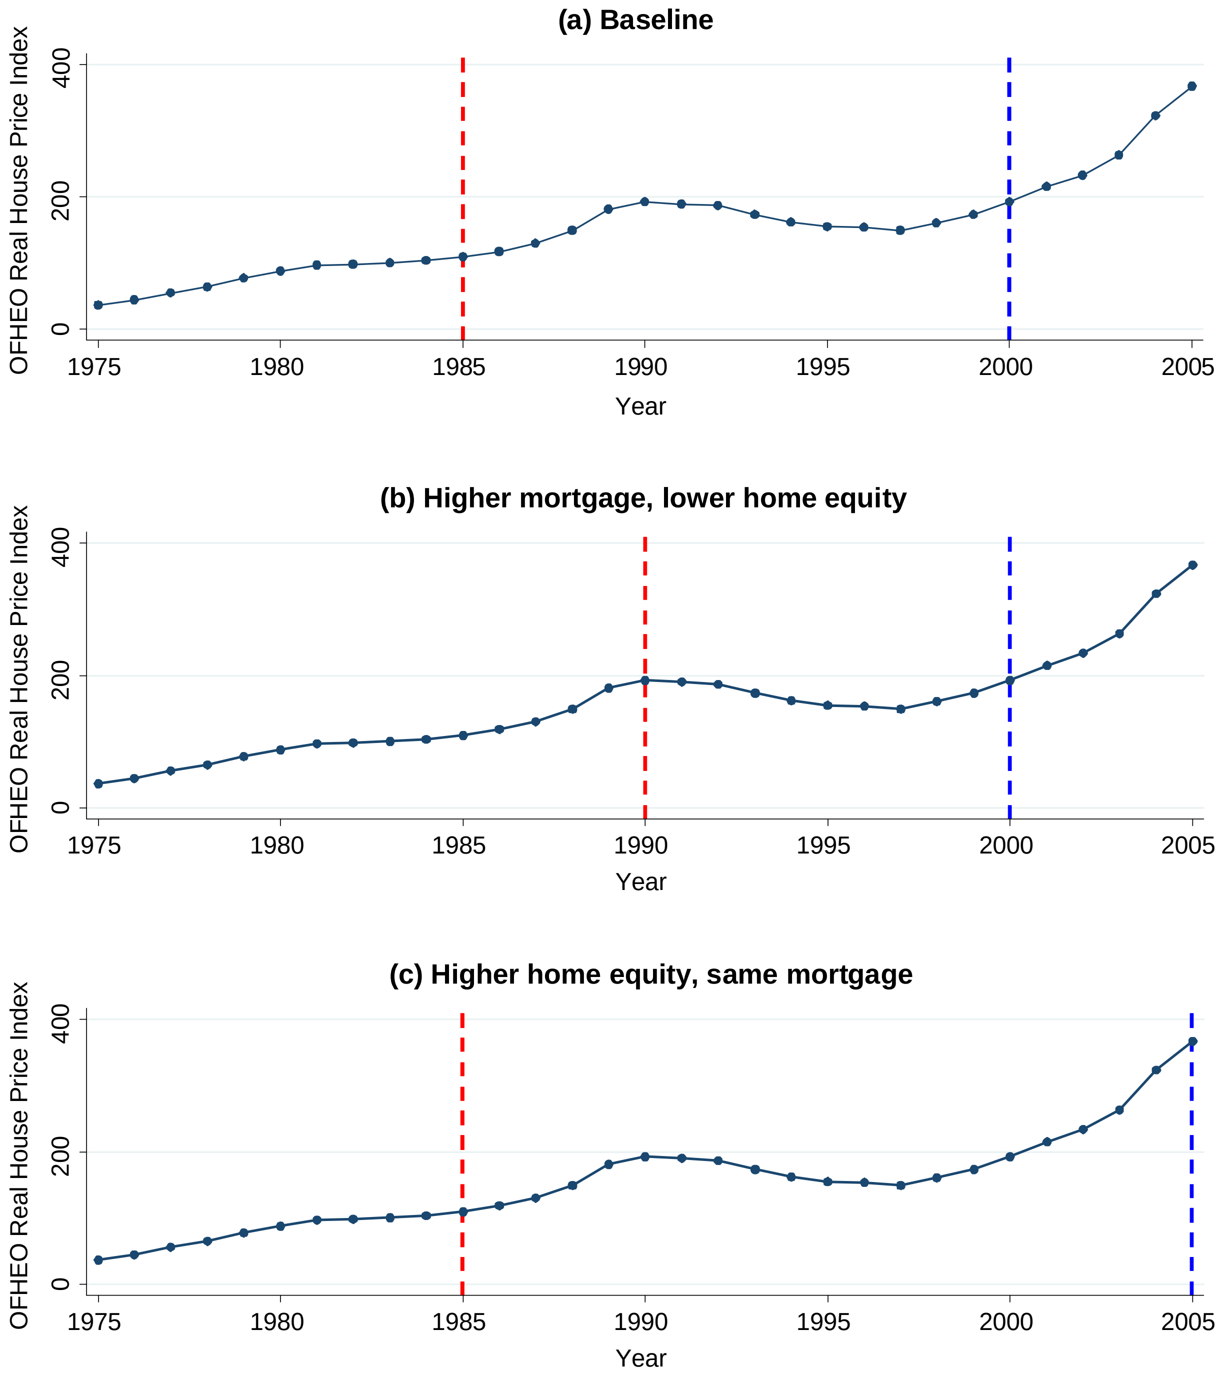

NOTE–This figure illustrates our identification strategy by showing the OFHEO real housing price index in California from 1975 to 2005. Panel A depicts an individual who buys a house in 1985 and whose portfolio is observed in 2000. Panel B shows an individual who buys the same house in 1990 instead of 1985, and has approximately $100,000 more mortgage debt when observed in 2000 as a result. Panel C shows an individual who buys in 1985 and is observed in 2005. This individual has approximately $175,000 more home equity than individual A.