

The Effect of Housing on Portfolio Choice*

Raj Chetty

Adam Szeidl

Harvard and NBER

Central European University, UC-Berkeley, NBER, CEPR

February 2012

Abstract

Economic theory predicts that home ownership should generally have a negative effect on risk-taking in financial portfolios, a result that affects the optimal design of a wide variety of financial and insurance policies. However, empirical work has not found a strong relationship between housing and portfolios. We identify two reasons for the divergence between the theory and data. First, it is critical to distinguish between home equity wealth and mortgage debt, as they have opposite-signed effects on portfolio choice. Second, it is important to isolate variation in home equity wealth and mortgage debt that is orthogonal to unobserved determinants of portfolios. We estimate a model that permits home equity wealth and mortgage debt to have different effects on portfolio shares. We isolate plausibly exogenous variation in home equity and mortgages by using differences across housing markets in average house prices and housing supply elasticities as instruments. Using data for 60,000 households, we find that increases in mortgage debt reduce stockholding significantly, while increases in home equity wealth raise stockholding. On net, an individual with 10% more mortgage debt and home equity has a 3% lower portfolio share of stocks. We conclude that housing has substantial impacts on portfolio choice, as theory predicts.

*E-mail addresses: chetty@fas.harvard.edu and szeidl@econ.berkeley.edu. Thanks to Thomas Davidoff, Edward Glaeser, Albert Saiz, Stephen Shore, anonymous referees, and numerous seminar participants for helpful comments. Gregory Bruich, Andras Komaromi, Jessica Laird, Keli Liu, James Mahon, Juan Carlos Suarez Serrato, and Philippe Wingender provided outstanding research assistance. We are grateful for funding from National Science Foundation Grants SES 0522073 and 0752835. Szeidl thanks the Alfred P. Sloan Foundation for support.

1 Introduction

Houses are the largest assets owned by most households, but the impact of housing on financial markets remains unclear. Theory predicts that housing generally reduces the demand for risky assets because it increases a household's exposure to risk and illiquidity (Grossman and Laroque 1990, Brueckner 1997, Flavin and Yamashita 2002, Chetty and Szeidl 2007). But empirical studies have not found a systematic relationship between housing and portfolios in practice (Fratantoni 1998, Heaton and Lucas 2000, Yamashita 2003, Cocco 2005).

This paper reconciles the theory with the data. We identify two key factors, one theoretical and one empirical, that explain the discrepancy. Theoretically, we show that it is critical to separate the effects of mortgage debt from the effects of home equity to characterize the effects of housing on portfolios. Empirically, we show that the endogeneity of housing choice biases prior estimates. Accounting for these two factors, we find that increases in mortgage debt induce substantial reductions in the share of liquid wealth held in stocks, while increases in home equity wealth raise stock ownership.

We structure our empirical analysis using an analytically tractable model of portfolio choice that incorporates both the illiquidity and price risk effects of housing. The model predicts that mortgage debt and home equity wealth have opposite-signed effects on portfolio choice. Increases in property value (holding wealth fixed) generally reduce the stock share of liquid wealth by increasing illiquidity, increasing exposure to risk, and reducing the present value of lifetime wealth. Increases in home equity raise the stock share of liquid wealth with CRRA preferences through a wealth effect, as emphasized by Yao and Zhang (2005). Because property value is the sum of mortgage debt and home equity, changes in property value conditional on the level of home equity are equivalent to changes in mortgage debt and should reduce stockholding. However, a regression of portfolio shares on property value that does not fully control for wealth – as in prior empirical studies – may yield ambiguous estimates because the variation in property values could be driven by an increase in home equity.

Because both portfolios and housing are endogenous choices that are affected by unobserved factors such as background risk (Campbell and Cocco 2003, Cocco 2005), one cannot identify the causal effect of housing on portfolios using cross-sectional variation across households. We use three research designs to address this central endogeneity problem. Each

strategy isolates variation in mortgage debt and home equity that is orthogonal to unobserved determinants of portfolio choice under a different set of assumptions.

Our first research design 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. 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.

We implement this cross-sectional IV strategy using microdata on housing and portfolios for 64,191 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 the share of stockholdings. A \$10,000 increase in property value (holding fixed home equity wealth) causes the stock share of liquid wealth to fall by \$310, or 5.5% of mean stockholdings in the sample. 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 5.9% through a wealth effect.¹ The elasticity of the stock share of liquid wealth with respect to outstanding mortgage debt is -0.3, while the elasticity with respect to home equity wealth is 0.4. These portfolio changes are driven by both the extensive and intensive margins: changes in mortgage debt and home equity wealth induces changes in both the probability of owning any stocks and the amount of stocks held conditional on stock ownership.

One potential concern with our first research design is that state-level house price fluctuations may be correlated with other factors such as local labor market conditions that

¹To facilitate comparison between samples with different rates of stock market participation and hence different mean stock shares of liquid wealth, we report results in percentages, rather than percentage points.

directly impact portfolio choice. Our second research design addresses this concern. Here, we instrument for property values and home equity using the current and year-of-purchase *national* average of house prices interacted with the *state* housing supply elasticity, as measured by Saiz (2010) based on land availability and regulations. Intuitively, fluctuations in the national housing market generate larger price fluctuations in states with inelastic housing supply, generating differential variation in house prices across states over time. This strategy yields estimates that are very similar to the first design. We estimate that a \$10,000 increase in property value causes a reduction in the stock share of liquid wealth of approximately 5.1%, while a \$10,000 increase in home equity (holding fixed total property value) raises the stock share by 4.7%.

Our third research design uses panel data to study the short-run dynamics of portfolios from the year before to the year after home purchase. We test whether individuals who buy a larger house reduce their stock share of liquid wealth more than those who buy smaller houses. We again 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 approaches 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 further mitigates concerns about the endogeneity of housing choices by permitting household fixed effects. Because the SIPP is a short panel, we observe portfolio shares both before and after home purchase for only 6,510 households. For this subset of households, we find that a \$10,000 increase in the price of the house leads to a 4.1% reduction in the stock share of liquid wealth in the year after home purchase, again very similar to the estimates from the first two designs. This finding shows that stockholders primarily sell stocks (rather than bonds) to finance down payments.

The magnitudes of the impact of housing on financial portfolios can be assessed by considering various counterfactuals. First, suppose households had the same level of home equity wealth but had their mortgage debt obligations cancelled. Our estimates from our the first research design imply that the stock share of portfolios would be 4.3 percentage points higher in this scenario. Given a mean level of liquid wealth in our sample of \$40,000 (in 1990 dollars), this translates into a \$1,700 increase in stockholdings per household in the sample. While this may appear to be a small change in absolute terms, it constitutes a 27% increase

in the stock share of liquid wealth relative to the sample mean. Second, suppose households had no mortgage debt and no home equity wealth. The net impact of having no housing wealth or liabilities would be an increase in the stock share of 5.2 percentage points (32%), or \$2,100. Finally, as another metric, a one standard deviation increase in mortgage debt reduces the stock share of liquid wealth by 4 percentage points (25%). This is similar to the impact of a one standard deviation increase in log financial wealth on stock shares (Calvet, Campbell, and Sodini 2007). The magnitudes of these impacts are also consistent with numerical simulations of our model when calibrated with parameters used in the literature.

Our estimates of the effect of housing on portfolios are larger and more robust than previous estimates. Fratantoni (1998) finds an elasticity of stock share with respect to mortgage debt of -0.15. In contrast, Heaton and Lucas (2000), Cocco (2005), and Yao and Zhang (2005) show that when 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. Yamashita uses age, family size, home tenure, and aggregate housing returns as instruments for mortgage debt; unfortunately, these instruments are unlikely to be valid because standard models (e.g. Cocco 2005) generate direct relationships between all of these variables and portfolio choice, independent of the housing channel. Consistent with these prior studies, we also find that OLS estimates in our data are often wrong-signed and are sensitive to covariates. Our IV estimates are less sensitive to specification because they are driven by variation that is orthogonal to most household-level determinants of portfolios and because we systematically separate the effects of mortgage debt and home equity.

The link between housing and financial decisions that we document here has implications for several issues. For example, our results suggest that recent increases in leverage due to the easing of credit in the U.S. (Mian and Sufi 2011) may have induced households to withdraw funds from the stock market. This reduction in demand for risky assets could have further precipitated the sharp decline in asset prices. Our results also support the view that homeownership amplifies the welfare cost of shocks. Policies such as unemployment and health insurance or restrictions on the riskiness of financial portfolios could therefore generate significant welfare gains for individuals who own houses or other risky, illiquid assets.

The remainder of the paper is organized as follows. The next section presents a portfolio choice model, analyzes its comparative statics with respect to housing, and quantifies the impacts one should expect using numerical simulations. Section 3 describes the data. Section 4 presents the empirical results. Section 5 concludes.

2 Theoretical Predictions

In this section, we characterize the forces through which exogenous changes in property value and home equity affect household portfolios. We begin by deriving an approximate analytical expression for optimal portfolio shares in a stylized two-period model. This stylized model provides a simple, tractable framework that unifies the intuitions of several papers that have highlighted different mechanisms through which housing affects portfolio choice, including illiquidity (Grossman and Laroque 1990, Chetty and Szeidl 2007), home price risk (Flavin and Yamashita 2002), hedging effects (Sinai and Souleles 2005), and diversification effects (Yao and Zhang 2005). We then generalize the model to allow for fixed moving costs, multiple periods, and labor income risk. Using numerical simulations, we show that the key comparative statics of the two-period model hold with plausible parametrizations in this more general environment.

2.1 Stylized Two-Period Model

Our stylized model builds on Cocco's (2005) model of housing and portfolio choice, which incorporates all of the mechanisms described above but does not permit an analytical solution. To obtain an analytic expression for portfolio shares, we make a number of simplifying assumptions, most importantly that households can only move at exogenous random dates. The more realistic model in which households can move by paying a fixed cost is analytically intractable, and we therefore analyze it using numerical simulations below.

A household endowed with a house H_0 , mortgage debt M_0 , and liquid wealth L_0 makes a financial portfolio investment decision in $t = 0$. Consumption takes place in $t = 1$, and the household maximizes

$$E_0 \frac{\left[C_1^{1-\mu} H_1^\mu \right]^{1-\gamma}}{1-\gamma} \tag{1}$$

where C_1 is adjustable (e.g., food) consumption and H_1 is housing consumption. As in

Campbell and Cocco (2003), we assume that moves in $t = 1$ are exogenous. With probability θ the household stays in the current house ($H_1 = H_0$), while with probability $1 - \theta$ it moves, and chooses H_1 optimally. One interpretation of this assumption is that the fixed cost of moving is sufficiently high that except for life-changing events, such as marriage or childbirth, the household does not consider changing houses when making financial investments. In this model, θ measures the strength of housing commitment.

At $t = 0$ the household can invest in a riskfree financial asset with return $1 + R_f = \exp(r_f)$ and a risky asset with return $1 + R = \exp(r)$, where r is normally distributed with mean μ and variance σ^2 . The only choice variable at $t = 0$ is α , the share of the risky asset out of liquid wealth. Let $R_p = \alpha R + (1 - \alpha) R_f$ denote the household's financial return, and assume that short sales constraints restrict $\alpha \in [0, 1]$. Home prices are $P_0 = 1$ and $P_1 = \exp(p_1)$, where p_1 is normal with mean μ_p and variance σ_p^2 . The correlation between home price growth and stock returns is $\rho = \text{corr}[p_1, r]$.

The household chooses α to maximize (1) subject to the budget constraint

$$C_1 + P_1 H_1 = (1 + R_p) L_0 + Y_1 + P_1 H_0 - (1 + R_m) M_0$$

where R_m is the mortgage rate and Y_1 is labor income, which for now we assume is deterministic. Let the present values of mortgage debt, labor income, liquid wealth, home value, and lifetime wealth be denoted by $M = M_0 (1 + R_m) / (1 + R_f)$, $Y = Y_1 / (1 + R_f)$, $L = L_0$, $PH = P_0 H_0 \cdot e^{\mu_p} / (1 + R_f)$, and $W = L + Y + PH - M$.²

Optimal portfolio shares. We derive an approximate equation for the optimal risk share α using log-linearization. Household optimization yields the following log-linear Euler equation:

$$Er - r_f + \frac{\sigma^2}{2} = \theta^* \cdot \text{cov}[r, -v'_{nm}] + (1 - \theta^*) \cdot \text{cov}[r, -v'_m],$$

where v'_{nm} and v'_m are the log marginal utilities of wealth in $t = 1$ in the “no move” and “move” states of the world and the weight

$$\theta^* = \frac{1}{1 + \frac{1-\theta}{\theta} \frac{\mu^{\mu(1-\gamma)}(1-\mu)^{-\mu-\gamma+\mu\gamma}}{(PH/W)^{\mu(1-\gamma)}(1-PH/W)^{-\mu-\gamma+\mu\gamma}}}. \quad (2)$$

²With the exception of housing, these measures are risk-adjusted present values. For housing, our definition does not account for house price risk but turns out to be suitable for the log-linearization approach used below.

Intuitively, the agent optimizes by trading off the expected gain from investing in the risky asset with the additional fluctuation in marginal utilities he bears as a result of the investment. The additional risk is measured by the covariance of the market return with marginal utilities, weighted by θ^* . The weight θ^* can be interpreted as a marginal-utility-adjusted probability of not moving, analogous to a state-price density. When the housing share of lifetime wealth PH/W equals the optimal share μ , $\theta^* = \theta$. But when $PH/W > \mu$, $\theta^* > \theta$: since the household starts with too much housing and too little adjustable consumption, the marginal utility of wealth is—on average—relatively higher in the no-move state. As a result, the consumer is more sensitive to fluctuations in this state, explaining the larger weight θ^* .

The optimal portfolio share can be derived from the Euler equation using standard methods (see e.g., Campbell and Viceira 2002):

Proposition 1 *Letting $\gamma^c = \mu + \gamma - \gamma\mu$, the optimal share of stocks out of liquid wealth at $t = 0$ is, to a log-linear approximation,*

$$\alpha = \frac{\mu - r_f + \sigma^2/2}{\sigma^2 \left[\theta^* \gamma^c \frac{L}{W-PH} + (1 - \theta^*) \gamma \frac{L}{W} \right]} + \text{cov}[p_1, r] \cdot (1 - \theta^*) \frac{\mu(\gamma - 1) - \gamma \frac{PH}{W}}{\sigma^2 \left[\theta^* \gamma^c \frac{L}{W-PH} + (1 - \theta^*) \gamma \frac{L}{W} \right]}. \quad (3)$$

Proof. See Appendix.

To understand this expression, first consider the case in which house prices do not covary with stock prices ($\text{cov}[p_1, r] = 0$). In this case, the second term drops out and (3) has an interpretation analogous to a familiar “myopic” rule: the numerator measures the expected excess return of stocks, while the denominator equals stock market risk σ^2 multiplied by effective risk aversion over liquid wealth. Because housing is a fixed commitment, risk aversion is the weighted average $\theta^* \gamma^c L / (W - PH) + (1 - \theta^*) \gamma \cdot (L/W)$. When the consumer is free to move ($\theta = \theta^* = 0$), this term simplifies to $\gamma L/W$, yielding the classic Merton (1969) formula adjusted for the fact that stocks are measured as a share of liquid rather than total wealth. When the consumer can never adjust housing ($\theta = \theta^* = 1$), effective risk aversion is $\gamma^c L / (W - PH)$. Effective risk aversion differs when the agent does not move for two reasons. First, because shocks are concentrated on adjustable consumption $W - PH$, they have an amplified effect on marginal utility (Chetty and Szeidl 2007). Second, because H_1 does not adjust, curvature is determined by $(1 - \mu)(1 - \gamma)$ in (1), generating the γ^c term.

When $\text{cov}[p_1, r] \neq 0$, home price risk generates a hedging demand for stocks, reflected in the second term in (3). This term is also affected by the strength of the housing commitment θ through θ^* . When $\theta = \theta^* = 1$, the home is never sold, and hence home price risk does not affect behavior (Sinai and Souleles 2005).

Comparative Statics. We are interested in characterizing how the optimal portfolio share varies with property value PH and total wealth W . With CRRA utility, the household seeks to maintain a constant share of its total wealth in risky assets as W rises. Therefore, an exogenous increase in home equity wealth – which is relatively safe – induces the household to buy more stocks. This “diversification effect” (Yao and Zhang 2005) is captured by the terms involving W in the denominator in (3). In our model, an increase in wealth also reduces θ^* , the weight of the no-move state in the Euler equation. Because the no-move state is typically riskier, this additional effect generally acts to further raise α^* .

Exogenous increases in property value PH reduce α^* through three channels. First, for a given W , increasing PH implies that a larger share of wealth is “tied up” in housing, making marginal utility higher and more sensitive to shocks in the no-move state. This effect arises from an increase in effective risk aversion $\gamma^c L / (W - PH)$ in the denominator of (3) and by a higher weight θ^* on the no-move state. Second, when $\text{cov}[p_1, r] > 0$, a higher PH results in greater exposure to home price risk, which has a negative effect on hedging demand. Third, holding fixed home equity, a higher property value means higher mortgage debt. If the mortgage rate exceeds the risk free rate ($R_m > R_f$), increased mortgage payments reduce lifetime wealth W , resulting in lower stockholdings in (3).

These comparative statics show that it is critical to distinguish changes in property value from changes in home equity wealth to uncover the effects of housing on portfolio choice. Increases in property value that come from more mortgage debt reduce stockholding, while increases in property value that are accompanied by additional home equity wealth have ambiguous effects.

2.2 Numerical Simulations

We now assess the quantitative importance of the comparative statics as well as their robustness to incorporating the following features: (1) fixed adjustment costs, which permit households to move at any time by paying a cost, (2) multi-period dynamics with persistent

uncertainty, (3) labor income risk, and (4) prepayment and refinancing of mortgages. Because these features make the model analytically intractable, we use numerical simulations to characterize the relationship between housing and portfolios.³

We begin by calibrating the parameters of the model based on the existing literature. For parameters related to life-cycle portfolio choice, we follow Cocco, Gomes, and Maenhout (2005) and set the annual riskfree return at $R_f = 0.02$, the annual equity premium at $ER - R_f = 0.04$, the standard deviation of the log stock return at $\sigma = 0.157$, and the coefficient of relative risk aversion at $\gamma = 10$. For housing related parameters, we set the relative preference for housing at $\mu = 0.2$ (Yao and Zhang 2005) and the annual standard deviation of home prices at $\sigma_p = 0.062$ and the annual mortgage rate at $R_m = 0.04$ (Cocco 2005). Cocco and Yao and Zhang assume a zero correlation between housing and the stock market; we report results both with $\rho = 0$ and with $\rho = 0.1$ here.

In our baseline case, we consider a ten year investment horizon, chosen to reflect the average age in our sample of 48 years. We also report results from simulations with a twenty year investment horizon. Cocco (2005) estimates a moving probability of 24.4% over a five year horizon, which implies that the probability a household does not move over a ten year horizon is about 57%. We therefore set a baseline $\theta = 0.55$, but evaluate the sensitivity of our results to other values as well. We assume that the household has liquid wealth $L_0 = \$40,000$, home value $P_0H_0 = \$125,000$ and mortgage $M_0 = \$53,000$, the sample means in our data. Cocco, Gomes, and Maenhout (2005, Figure 3b) report that the ratio of the risk-adjusted present value of future labor income to current financial wealth is about 5 for households in their late forties and early fifties. We therefore set labor income at $Y_1 = 5L_0$.

Simulations of stylized model. We first report numerical simulations of the stylized two-period model as a reference to verify that the approximate solution in (3) accurately captures the comparative statics of the model.⁴ Table 1 reports optimal portfolios as a function of property value and home equity for a range of model parameters. Panel A confirms that increases in property value P_0H_0 (holding fixed home equity wealth) reduce the optimal share

³Details on the simulations methods are given in the appendix. We thank Joao Cocco for sharing his code for solving the model in Cocco, Gomes and Maenhout (2005).

⁴The quality of the approximation is quite high for short horizon but deteriorates over longer horizons. For example, the average absolute difference between the numerical and the approximate solution across all the parameter values considered in Table 1 is only 0.55 percentage points for a one year horizon. With a five year horizon, the mean error grows to 2.07 percentage points and for ten years it is 2.98 percentage points. Despite these deviations, the approximate solution shows the same patterns as the numerical results.

of stocks. For example, when the probability that the household does not move is $\theta = 0.55$ and housing is uncorrelated with returns ($\rho = 0$), increasing property value from \$125,000 to \$135,000 results in a reduction in stock share from 66.5% to 60.1%, or by about 9.6%.⁵ Panel B considers the effect of changes in home equity wealth. In the baseline case, an increase in home equity from \$72,000 to \$82,000, while holding home value fixed at \$125,000, increases the stock share from 66.5% to 72.7%, or by about 9.3%. We observe qualitatively similar effects for other parameter values. Having established the comparative statics of interest in our stylized model, we now consider a series of generalizations.

Fixed adjustment costs. We begin by relaxing the assumption that households can only move at random, exogenous dates. A more realistic assumption is that households can move at any time by paying a fixed cost. Let λ denote the size of this fixed cost as a share of property value. Smith, Rosen, and Fallis (1988) estimate the monetary component of moving costs to be approximately $\lambda = 0.1$. We consider values of $\lambda = 0.1$ and $\lambda = 0.2$, the latter of which captures other utility costs of moves (e.g., the need to change a child’s school). Panel A of Table 2 reports simulations analogous to those in Table 1 from this model. The direction of comparative statics are generally the same, but the property value effects are smaller in magnitude, as should be expected given that housing is a weaker commitment in this model.

In the fixed cost model, the comparative statics of interest actually change sign for some parameter values. For instance, when $\lambda = 0.2$, increasing home value from \$105,000 to \$115,000 reduces the stock share from 84.9% to 82.6%, but an additional increase in home value to \$125,000 actually *increases* the stock share slightly, to 82.8%. Such non-monotonicities in risk preferences in the presence of fixed costs were first observed by Grossman and Laroque (1990) and more extensively documented by Yao and Zhang (2005). The intuition is that households who are relatively close to the boundary of their inaction region have a gambling motive: by holding more stocks, they can increase the probability of buying their “ideal” house. For households who are on the margin of moving, this mechanism can

⁵Our calibrated model produces stock shares that are substantially higher than the mean stock share of liquid wealth in our data (16%). The model matches the data better if we focus on the subsample of stockholders, among whom the average stock share is 55%. As Cocco, Gomes and Maenhout (2005) and Davis, Kubler and Willen (2006) emphasize, calibrated models of portfolio choice frequently overpredict stockholdings because relatively safe future labor income creates an incentive to leverage financial investment. The discrepancy between the model and the data in levels could be addressed by allowing for fixed costs of participating in asset markets, which make a large subset of households choose not to participate (Vissing-Jorgensen 2002).

sometimes overpower the three forces that act toward reducing α^* . However, Table 2 shows that for most parameter values, the other three forces dominate. Therefore, given a plausible distribution of wealth and housing values in the population, we expect the average effect of property value on the stock share of liquid wealth to be negative.

Persistent uncertainty. Our stylized model is effectively static, because all uncertainty is resolved in a single period. To consider the effects of persistent uncertainty, we now introduce a third period into the baseline model. In $t = 0$, the household makes a portfolio decision as above. In $t = 1$, it repays its outstanding mortgage, earns labor income Y_1 , moves houses with (exogenous) probability $(1 - \theta)$, consumes, and makes a new portfolio decision. In $t = 2$, the household earns labor income Y_2 , moves with independent probability $(1 - \theta)$, and consumes. We assume that each period lasts ten years. We set $Y_1 = 4L_0$, $Y_2 = 3L_0$, and set the remaining parameters to the value above, with an annual discount factor of $\delta = 0.96$.⁶ Panel B of Table 2 shows that in this environment with persistent uncertainty, the effects of property value and home equity are similar to the baseline specification, but larger in magnitude.

Income risk. A third simplification in our baseline model is the absence of labor income risk. To account for income risk, we follow Cocco, Gomes and Maenhout (2005) and introduce a shock which occurs with probability of 5% per period and results in a loss of 20% of labor income.⁷ As Panel C of Table 2 shows, introducing these shocks reduces the stock share of portfolios. This is because these shocks reduce expected lifetime wealth and because they increase risk aversion. However, the predictions about the effect of property value and home equity are again unchanged in sign and remain relatively similar in magnitude.

Prepayment and refinancing. A final simplification in the baseline model is that mortgage payment 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 emphasized by Campbell and Cocco (2003) is that borrowing costs, such as costs of refinance, generate a motive to hold a buffer stock in liquid assets. One simple way to model such costs of short term borrowing is to assume that adjustable consumption C_1 must

⁶We introduce a discount factor only here because in the baseline model, utility is obtained in a single period, and hence discounting has no effect on behavior.

⁷One can derive an approximate analytic formula for this case with income risk, but in the interest of space we do not pursue this approach here.

be financed from a share η of labor income Y_1 . Here ηY_1 can be interpreted as labor income which accumulates before consumption, or which can be borrowed against. To see the effect of this constraint, note that in the no move state, optimal consumption is $C_1^* = L_1 + Y_1 - M_1$, while borrowing limits impose $C_1 \leq L_1 + \eta Y_1$. Thus, whenever $M_1 < (1 - \eta) Y_1$ the constraint is binding. In this case, holding mortgage debt is useful because it substitutes for borrowing against future labor income. Using this insight, it can be shown that for a range of $\eta < 1$ and $\theta > 0$, optimal mortgage debt is $M_0 = (1 - \eta) Y_1 / (1 + R_m) > 0$ and that, given this M_0 , the optimal portfolio coincides with the solution of our baseline model.

Across the model specifications, the effect of a \$10,000 increase in property value ranges from an average reduction in the stock share of about 2.5 percent in the fixed adjustment cost model to an average reduction of about 20 percent in the model with exogenous random moves and persistent uncertainty. The average effect of a \$10,000 increase in home equity wealth ranges from an increase of 5-23 percent in stock shares. While the simulations show that mortgage debt generally reduces risk taking in financial portfolios and home equity wealth increases it, the magnitudes of these effects are sensitive to model specification. The quantitative impacts of housing on portfolio choice are therefore an empirical question, to which we now turn.

3 Data and Sample Definition

We estimate equation (4) using data from seven Survey of Income and Program Participation panels that began in years 1990-2001. Each SIPP panel tracks 20,000 to 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 Survey of Consumer Finances are its large sample size and detailed information about covariates such as a complete housing history and geographic location.

We obtain quarterly data on average of housing prices by state from 1975-2004 using the repeat sales index constructed by the Office of Federal Housing Enterprise Oversight (OFHEO). Calhoun (1996) provides a detailed description of the construction of the OFHEO index, which has been widely used in studies of housing markets (see e.g., Himmelberg, Mayer,

and Sinai 2005).⁸ We obtain land topology-based measures of housing supply elasticity by state from Saiz (2010, Table 6). Saiz predicts housing supply elasticities using data on physical and regulatory constraints (land availability and use regulations), providing a convenient index of the supply constraints in each housing market.

The seven SIPP panels together contain information on 163,405 unique households, of which 97,798 are homeowners.⁹ 70,924 of these households bought their current house after 1975 and therefore have OFHEO data for the year of home purchase, which is required for our instrumental variable analysis. We exclude an additional 6,733 households whose reported liquid wealth by our definition is zero, making their portfolio shares ill defined. These exclusions leave us with 64,191 homeowners in our cross-sectional analysis sample.

Table 3a reports summary statistics for the cross-sectional analysis sample. In the cross-sectional sample, homeowners own houses that are worth approximately \$125,000 on average in 1990 dollars. The average amount of home equity is \$72,000 and the average outstanding mortgage is \$53,000. The average household head is 48 years old and has lived in his current house for 8.4 years. Mean total wealth (which includes liquid wealth, home equity, and other illiquid assets such as cars) is \$173,000.¹⁰ We define liquid wealth as the sum of assets held in stocks, bonds, checking, and savings accounts (excluding retirement accounts). Mean liquid wealth is \$40,000, but this distribution is very skewed; the median level of liquid wealth is only \$5,600.¹¹ Households hold on average approximately 16% of their liquid wealth in the form of stocks in taxable (non-retirement) accounts and 84% in “safe” assets (bonds, checking, and savings accounts). The relatively small fraction of wealth held in stocks reflects the fact that only 29% of the households in the data hold stocks outside their retirement accounts, consistent with Vissing-Jorgensen (2002).

⁸We use OFHEO price indices rather than other popular measures such as Case-Shiller indices because Case-Shiller data are available only starting in 2000 for selected metro areas. Unfortunately, geographic information below the state level is not available for more than two-thirds of the observations in our sample. Although the two indices differ in the way they treat appraisals and the set of loans they consider, Leventis (2007) reports a correlation of 0.98 between the OFHEO and Case-Shiller indices for markets where both measures are available.

⁹See Appendix Table 1 for summary statistics for the full SIPP sample.

¹⁰Total wealth is gross household wealth measured on the survey. It includes financial assets as well as all real estate (including second homes), cars, and private business equity. Debts are not subtracted from the total wealth measure.

¹¹Skewness and outliers do not affect the results reported below. Trimming outliers (e.g. by excluding the top and bottom 1% of households by wealth or property value) has virtually no effect on our 2SLS estimates. This is because the distribution of predicted housing values generated by the instruments is not skewed. There are few outliers in the fitted values from the first stage.

Panel data on portfolio shares are available for households in the 1996 and 2001 SIPP panels. In these panels, data on portfolio shares were collected annually, giving us information on assets and homeownership between 3 to 4 times per household. We form our panel analysis sample using the 6,150 observations for which we observe a purchase of a new house within the panel and have data on portfolio shares both before and after this home purchase.¹²

Table 3b reports summary statistics for the sample we use in the panel analysis. 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 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.

4 Empirical Analysis

We estimate the impacts of property value and home equity using the following linear specification for portfolio shares:

$$\text{stock share}_i = \text{const} + \beta_1 \text{property value}_i + \beta_2 \text{home equity}_i + \gamma X_i + \varepsilon_i \quad (4)$$

where X_i denotes a vector of controls, including components of total wealth such as liquid wealth and income. The model in Section 2 predicts $\beta_1 < 0$ and $\beta_2 > 0$.¹³ 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 in income (Cocco 2005).

Some of the effects captured by the error term may be correlated with property value, creating bias in OLS estimates of β_1 and β_2 . For instance, Cocco (2005) emphasizes biases due to unobserved labor income, which affects both the stock share and property value.

¹²When 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.

¹³An alternative specification is to normalize the housing variables by liquid wealth. We show that our results are robust to such a specification, but opt to use levels in our baseline model for two reasons. First, when liquid wealth is imperfectly measured and close to zero for some observations, normalizing by it introduces large outliers in the independent variables of interest. Second, our simulations show that one should find a relationship between the stock share and levels of property value and home equity wealth.

Suppose that $Y_1 = Y_1^{\text{obs}} + Y_1^{\text{un}}$ where only Y_1^{obs} is observed to the econometrician. Since higher lifetime wealth generates higher stockholdings, ε is positively related to Y^{un} . If households with higher future labor income own larger houses – as predicted by the model with persistent uncertainty – property value is also positively related to Y^{un} , and hence the OLS estimate of β_1 is biased upward. Indeed, Cocco (2005, Table 6) runs cross-sectional OLS regressions using simulated data from his model and finds a *positive* effect of mortgage debt on stockholdings caused by omitting future labor income from the regression. Such endogeneity problems make it essential to isolate variation in property value and home equity that is orthogonal to ε in order to identify β_1 and β_2 .

We divide our empirical analysis into four sections. First, we confirm that estimating (4) using OLS in our data yields results that are similar to those of prior studies. We then identify the causal impacts of mortgage debt and home equity wealth on portfolios by using three different research designs to estimate (4): variation in mean house prices, variation in local housing supply constraints, and changes in portfolio shares around home purchase in panel data.

4.1 OLS Estimates

Previous studies have estimated OLS regressions of portfolio shares on property values, mortgage debt, and home equity with various control vectors and obtained mixed results. To ensure that the differences between our findings and theirs are not driven by differences in data or sample definitions, we begin by estimating similar specifications in our sample.

Column 1 of Table 4 reports OLS estimates of a regression of the stock share of liquid wealth on property value and home equity wealth without any covariates. Consistent with the findings of Heaton and Lucas (2000), Cocco (2005), and Yao and Zhang (2005), we find that an increase in property value (mortgage debt) is *positively* associated with the stock share of liquid wealth, contrary to the model’s predictions. This is presumably because individuals with larger properties tend to be wealthier or face less background risk and these omitted factors induce them to hold more stocks.

In column 2, we attempt to account for some of these factors by including controls for household income and private business wealth; household head’s education, number of children, and age; and a 10 piece linear spline in liquid wealth to control flexibility for a house-

hold's level of wealth. The inclusion of these covariates reduces the coefficient on property value by approximately 80%, but it remains positive in sign.

In column 3, we exclude households with zero mortgage debt, who constitute 23% of homeowners in the sample, as in Fratantoni (1998). This change in sample specification makes the coefficient on property value negative and statistically significant, consistent with Fratantoni's findings. Importantly, Fratantoni is not able to control for location as the SCF does not contain geographic information. Once indicators for state of residence are included, the negative correlation between property values and stock shares is no longer significant, as shown in Column 4 of Table 4.

These OLS results echo the instability of estimates found in prior studies. Moreover, they indicate that the endogeneity of housing choices is likely to bias the effect of property value on stock shares upward. These findings call for research designs that isolate variation in mortgage debt and home equity that is less correlated with unobserved determinants of portfolios.

4.2 Research Design 1: Mean House Prices

Identification Strategy. Our first research design 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 (4) purely from within-state changes in house prices.¹⁴ 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.

The first three columns of Table 5 report first stage regressions of mortgage, home equity, and property value (mortgage plus home equity) on the two instruments. These specifications include state, year of purchase, current year, and age fixed effects as covariates.¹⁵ These first-stage effects remain very similar when we include the following vector of “full controls,” which we use to evaluate robustness of each of our specifications below: household income, household head’s education, number of children, the state unemployment rate in the current year, private business wealth, and a ten piece linear spline for liquid wealth.

Column 1 of Table 5 shows that higher current house prices strongly predict higher property values, with a t-statistic of 40. Conditioning on current prices, higher house prices at the time of purchase predict slightly lower current property values, confirming that individ-

¹⁴Since the instruments only vary by state and year/year of purchase, one cannot include state×year effects (i.e. allow state effects to vary over time).

¹⁵By including both year of purchase and current year fixed effects, we control non-parametrically for home tenure.

uals 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.¹⁶

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 state housing markets could be correlated with fluctuations in the local labor market or other economic conditions, which might in turn directly influence portfolio choices. Second, the exclusion restriction could be violated via selection effects. People who buy houses when prices in their state 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 concerns below after presenting the baseline results.

2SLS Results. Columns 4-9 of Table 5 report two-stage least squares estimates of β_1 and β_2 in (4), where home equity and property value are instrumented using the two OFHEO price indices. In column 4, 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 \$10,000 increase in an individual’s mortgage debt reduces his stock share of liquid wealth by 0.89 percentage points (\$350). Given a mean stock share in the analysis sample of 16.1%, this is equivalent to a 5.5% reduction in the stock share of liquid wealth. This impact lies within the range of estimates implied by the numerical simulations in Tables 1 and 2. The point estimate is closest to the prediction of the fixed adjustment cost specification, which is perhaps the most plausible model. The elasticity of

¹⁶Refinancing does not affect our 2SLS estimates because it rescales both the first-stage and reduced-form coefficients by the same amount. Refinancing could affect liquid wealth; we account for this channel by conditioning on liquid wealth using a flexible spline in many of our specifications. Note that in a heterogeneous population, our IV strategy will estimate a local average treatment effect that applies to individuals who do not fully refinance their mortgages when property values go up.

the stock share of liquid wealth with respect to mortgage debt is approximately -0.3 at the sample mean mortgage debt of \$53,000.

The estimate of the home equity coefficient in column 4 implies that a \$10,000 increase in home equity raises the stock share by 0.95 percentage points (5.9%) when total property value is held fixed, which we interpret as a wealth effect. The mean home equity in the sample is approximately \$72,000, implying an elasticity of stock share of liquid wealth with respect to home equity wealth of approximately 0.4.

Column 5 of Table 5 replicates column 4 with the fixed effects and the full set of covariates: liquid wealth spline, private business wealth, education, income, number of children, and the state unemployment rate. The estimate of the property value coefficient is virtually unaffected by controls, unlike the OLS estimates in Table 4. Since controlling for observables has little impact on the estimate, one can be more confident that biases due to omitted unobservables are not driving these IV results.

To further interpret the magnitude of these effects, consider the following counterfactuals. First, suppose households had the same level of home equity wealth but had their mortgage debt obligations cancelled. The estimates from column 5 imply that the stock share of portfolios would be $0.52890 \times 7.78 = 4.1$ percentage points higher in this scenario. This is a 27% increase relative to the mean stock share of 16.1%. Second, suppose households had no mortgage debt and no home equity wealth. The net impact of having no housing wealth or liabilities on the stock share of liquid wealth would be $-7.78 \times 1.25154 + 6.22 \times 0.72301 = -5.23$ percentage points, a 32% reduction.

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. It follows 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.

Robustness Checks. In columns 6-8 of Table 5, we evaluate the robustness of our estimates to alternative specifications and sample definitions. In column 6, we estimate a model

analogous to column 4 using logs instead of levels for the independent variables. We instrument for $\log(\text{property value})$ and $\log(\text{home equity})$ 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 previous results, the estimates reveal that increases in property value significantly reduce the share of stocks in liquid wealth, and increases in home equity wealth increase stock shares.

Column 7 reports estimates from a specification analogous to column 5 except the endogenous regressors are also defined as shares of liquid wealth, 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 column 4 as instruments for these ratios. This specification effectively tests whether households with a large amount of mortgage debt to liquid wealth hold safer portfolios using a different functional form to account for variation in wealth. One problem with this specification is that we introduce substantial outliers, as there are many observations with near-zero liquid wealth. To reduce noise from these outliers, we exclude observations with ratios of property value or home equity to liquid wealth above 20. The estimates are consistent with those obtained previously, but less precisely estimated because of the instability of the ratios.

In column 8, we replicate the levels specification with the controls in column 5, 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 and home equity coefficients are slightly larger in magnitude than those in the full sample. Housing remains an important determinant of portfolio choice even for wealthier households.

Extensive and Intensive Margin Response. In columns 9 and 10, we decompose the effects of housing on stock shares into stock market participation decisions (whether to own any stocks) and intensive margin changes in portfolio allocations (how much money to invest in stocks conditional on owning stocks). Column 9 replicates column 5, replacing the dependent variable with an indicator for owning stocks. A \$10,000 increase in an individual's mortgage is estimated to reduce his probability of owning stocks by 1.4 percentage points, relative to a mean of 29%. Conversely, increases in home equity wealth increase the probability of stock

market participation by a similar magnitude.

Column 10 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 0.¹⁷ The estimates imply that a \$10,000 increase in mortgage debt reduces stock shares for stock market participants by 3.1 percentage points relative to a base of 55%. Home equity changes again have similar effects in the opposite direction.

Threats to Identification. 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 2SLS results in Tables 3. 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.

The first threat to a causal interpretation of the two reduced-form relationships is that fluctuations in current home prices are correlated with portfolios through omitted variables. For instance, house prices may be related to local economic conditions that directly affect portfolio choice. We believe that such effects are unlikely to be responsible for our findings for two reasons. First, controlling for observable measures of the local business cycle by using state unemployment rates and current household income has little effect on the estimates. Second, any remaining omitted variables (e.g. expectations of future labor income) are likely to bias the estimated effect of current house prices on stock shares upward. If individuals are unobservably wealthier when house prices are high in their area, their stock shares should rise because higher income individuals tend to hold more stocks. This would work *against* our

¹⁷Estimating a 2SLS model only on the subsample of stock market participants yields biased estimates because changes in home equity and mortgages affect stock market participation rates, generating selection effects.

second reduced-form result that fluctuations in property value have no effect on portfolios.

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 houses when house prices are relatively high may have different risk preferences. We believe that such selection bias is modest in our setting for the same two reasons. Controlling for observables has little impact on the estimates, indicating that selection on observables is minimal. And again, we expect such selection biases to work against our findings: those who are willing to buy a house when prices are relatively high are presumably *less* risk averse (Shore and Sinai 2010). This would work against our first reduced form finding that individuals who buy when prices are high (and thus have more mortgage debt) have *safer* portfolios.

While these arguments suggest that the results in Table 5 are unlikely to be driven by omitted variable and selection biases, they are not conclusive. We therefore turn to two alternative identification strategies that directly overcome these problems using different sources of variation.

4.3 Research Design 2: Variation in Housing Supply Elasticities

Identification Strategy. Our second research design exploits national house prices interacted with local housing supply elasticities to generate variation in home equity and property value. To understand the intuition for this strategy, consider two states, one with an inelastic housing supply (e.g. California) and another with highly elastic housing supply (e.g. Kansas). When there is an aggregate demand shock for housing at the national level, there is very little adjustment in the supply of housing in California, so prices covary strongly with the national prices. However, in Kansas, most of the adjustment takes place on the supply margin and local house prices are much more stable. More generally, aggregate demand shocks for housing (which we measure using national house prices) have larger impacts on house prices in states with low housing supply elasticities, generating differential variation in house prices across states (Mian and Sufi 2011). The advantage of this source of variation is that it avoids the potential for omitted variable bias due to local economic shocks because the variation is driven purely by national demand shocks.

To implement this strategy, we instrument for mortgage debt and home equity with

current and year-of-purchase *national* house prices interacted with the *state* housing supply elasticity. The housing supply elasticity is taken from Saiz (2010), who constructs predicted elasticities using measures of local physical and regulatory constraints. Including year fixed effects and state fixed effects in the regressions absorbs the level effects of the national prices and the state housing supply elasticity. Therefore, our instruments are simply the two interaction effects.

Columns 1-3 of Table 6 report first stage regressions of property value, home equity, and mortgage debt on these two instruments. The specifications in columns 1-3 include state, year of purchase, current year, and age fixed effects as covariates. As above, the first-stage estimates are unaffected by the inclusion of additional controls. Column 1 shows that higher national prices in the current year have significantly smaller effects on property values in states with highly elastic housing supply. Column 2 shows that higher national prices also have smaller effects on home equity in more elastic housing markets, as expected. Higher national prices at the time of purchase reduce home equity by a smaller amount in areas with elastic housing supply. Conversely, column 3 shows that higher prices at the time of purchase have smaller impacts on mortgage debt in elastic markets. All of these first-stage effects are highly significant, although the t statistics are somewhat smaller than in the first research design because this strategy exploits a narrower source of variation.

2SLS Results. Columns 4-6 of Table 6 report two-stage least squares estimates of β_1 and β_2 in (4), where home equity and property value are instrumented using the two national price indices interacted with state housing supply elasticity. In column 4, we estimate the model including current year, year of purchase, age, and state fixed effects. The point estimate of the property value coefficient implies that a \$10,000 increase in an individual's mortgage reduces his stock share of liquid wealth by 1.2 percentage points. A \$10,000 increase in home equity increases his stock share by 1.4 percentage points. Both estimates are statistically significant with $p < 0.01$. Column 5 of Table 6 replicates column 3 with the full set of covariates in addition to the fixed effects. Including the full set of controls does not have a statistically significant impact on the coefficient estimates. The magnitudes of the coefficients are quite similar to the corresponding coefficients in column 5 of Table 5, although slightly less precisely estimated because the first-stage has less power. Column 6 replicates column 5, replacing the dependent variable with an indicator for owning stocks to estimate the extensive

margin response. A \$10,000 increase in an individual’s mortgage is estimated to reduce his probability of participating in the stock market by 1.6 percentage points. Increases in home equity wealth increase the probability of stock market participation by a slightly smaller magnitude. These estimates are again fairly similar to those in Table 5, giving us greater confidence that our estimates are not significantly biased by omitted variables.

4.4 Research Design 3: Portfolio Changes Around Home Purchase

Identification Strategy. Our third identification strategy directly addresses concerns about selection by examining changes in portfolio shares within a household. Do individuals who buy more expensive houses reduce their stockholdings by a larger amount from the year before to the year after home purchase? We answer this question using the small subsample of households for whom we (1) observe a home purchase within our data and (2) observe portfolio shares both before and after home purchase. Note that this sample includes both individuals who transition from renting to owning and individuals who bought a new house within our sample frame. As discussed in Section 3, this panel sample includes much fewer households than the cross-sectional analysis sample because the SIPP tracks households for only 3 years and relatively few households buy a house within that window.

Define $\Delta x = x_{t+1} - x_{t-1}$ for an individual who buys a new house in year t . We estimate (4) in first differences:

$$\Delta \text{stock share}_i = \alpha + \beta_1 \Delta \text{property value}_i + \beta_2 \Delta \text{total wealth}_i + \gamma \Delta X_i + \Delta \varepsilon_i \quad (5)$$

This estimation strategy complements the preceding research designs by addressing selection directly. If our results are driven by selection effects, individuals who buy houses when prices are high would hold more conservative portfolios even *before* they buy their houses and we would not find $\beta_1 < 0$ in (5).

To account for the endogeneity of the size of the house one purchases, we instrument for $\Delta \text{property value}$ using the state house price index in the year of home purchase. Columns 1-3 of Table 7 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 a control for the change in total wealth. To reduce the influence of

outliers, we exclude 62 households who report changes in total wealth of more than 1 million dollars in these specifications; we show below that this exclusion has no effect on our estimate of β_1 but does affect the estimated wealth effects. The estimates show that individuals who buy houses in higher priced markets spend more on their houses. Most of the increase comes from taking on more mortgage debt rather than making a bigger downpayment to build home equity.

Because we only observe portfolio shares over two to three years, 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 preceding cross-sectional specifications. Instead, we control for the change in total wealth in (5) directly. To the extent that this approach fails to adjust adequately for the impacts of changes in wealth, our estimate of β_1 in the panel design will be biased toward zero because it captures not only the impacts of having more mortgage debt but also the impacts of having more wealth.

In practice, we find that controlling for the change in wealth has little impact on our estimate of β_1 because local house prices are not strongly correlated with changes in total wealth from the year before to the year after purchase. Intuitively, an individual who buys a house in more expensive market ends up with less liquid wealth but similar total wealth after the house purchase. As a result, the IV estimate of β_1 in (5) is effectively identified from changes in property value that are orthogonal to changes in total wealth. Therefore, we expect the estimates of β_1 from this design to be fairly comparable to the cross-sectional estimates of the impacts of mortgage debt on stockholding.

2SLS Results. Columns 4-6 report 2SLS estimates of the effect of changes in property value on the stock share of liquid wealth. In column 4, we include state, age, and year fixed effects and the change in total wealth as controls. A \$10,000 increase in property value is estimated to reduce the stock share by 0.9 percentage points in this specification. This estimate is statistically significant with $p < 0.01$. A \$10,000 increase in wealth is estimated to increase stock shares by 0.6 percentage points. Reassuringly, this estimate is quite similar to the estimated impacts of home equity wealth on the stock share from our first two identification strategies.

Column 5 shows that controlling for education, number of children, state unemployment

rate, and the change in household income does not affect the results significantly. Column 6 shows that the estimated impact of changes in property value on the stock share of liquid wealth remains unchanged when the outliers with wealth changes of more than 1 million dollars are included. Not surprisingly, however, these outliers substantially attenuate the estimated effect of wealth on portfolio shares. Finally, column 7 replicates column 5, replacing the dependent variable with an indicator for owning stocks. A \$10,000 increase in an individual's mortgage is estimated to reduce his probability of owning stocks by 0.7 percentage points relative to a baseline stock ownership rate of 38% in this sample.

The panel analysis confirms that the difference in portfolios between individuals who buy when house prices are high and low emerges immediately *after* home purchase, directly addressing concerns about selection bias. The similarity of the estimates from the three research designs indicates that mortgage debt has a robust negative effect on risk taking in financial portfolios over both short and long horizons.

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 0.4. On net, our estimates imply that stock shares of liquid wealth would be 5 percentage points (\$2,000 or 32%) higher in an economy without housing (no mortgage debt and no home equity wealth).

Our empirical results suggest 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 the stock market. 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 is consistent with the hypothesis that the illiquidity of housing amplifies household risk aversion. An interesting avenue for future research would be to explore whether fluctuations in the liquidity of housing markets over time induce changes in financial portfolios. It is also important to analyze whether the commitment of having to make mortgage payments – a “cash commitment” that arises from liquidity constraints – or the commitment of being unable to adjust housing consumption easily is what amplifies risk aversion.¹⁸ Depending upon which mechanism is more important, reducing transaction costs in the housing and mortgage markets could raise welfare both directly and by allowing households to bear more risk in their financial portfolios.

¹⁸Such an analysis would require variation in mortgage payments that is orthogonal to property value, perhaps arising from differences in the term structure of loans.

Appendix

Proof of Proposition 1. Let $W_1 = L_0(1 + R_p) + P_1H_0 - M_1 + Y_1$. If the household can move, the optimal consumption bundle satisfies $C_1 = (1 - \mu)W_1$ and $P_1H_1 = \mu W_1$, implying that utility in this state is

$$V_m(W_1) = \frac{[\mu^\mu (1 - \mu)^{1-\mu}]^{1-\gamma}}{1 - \gamma} \cdot \left(\frac{W_1}{P_1^\mu}\right)^{1-\gamma}. \quad (6)$$

If the household cannot move, the consumption bundle is $C_1 = W_1 - P_1H_0$ and $H_1 = H_0$, and hence utility is

$$V_{nm}(W_1) = \frac{H_1^{\mu(1-\gamma)}}{1 - \gamma} \cdot (W_1 - P_1H_0)^{(1-\mu)(1-\gamma)}. \quad (7)$$

We then define $V(W) = V_m(W)$ if the household moves and $V_{nm}(W)$ otherwise. The first order condition of the problem implies

$$E[(R - R_f) \cdot V'(W_1)] = 0.$$

We can write this as

$$\begin{aligned} & \theta \cdot E[(1 + R) \cdot V'_{nm}(W_1)] + (1 - \theta) \cdot E[(1 + R) \cdot V'_m(W_1)] \\ = & \theta \cdot E[(1 + R_f) \cdot V'_{nm}(W_1)] + (1 - \theta) \cdot E[(1 + R_f) \cdot V'_m(W_1)] \end{aligned}$$

where both sides are positive. To log-linearize each side separately, first introduce the notation that V_{nm}^{0r} and V_m^{0r} are the marginal utilities in the two states assuming that the agent has a fully safe financial portfolio, and $P_1 = \exp(\mu_p)$ without risk. Now take logs of the left hand side and use a first order approximation as follows:

$$\begin{aligned} & \log[\theta \cdot E[(1 + R) \cdot V'_{nm}(W_1)] + (1 - \theta) \cdot E[(1 + R) \cdot V'_m(W_1)]] \\ \approx & k + \theta^* [\log E[(1 + R) \cdot V'_{nm}(W_1)]] - \log[(1 + R_f) V_{nm}^{0r}] \\ & + (1 - \theta^*) [\log E[(1 + R) \cdot V'_m(W_1)]] - \log[(1 + R_f) V_m^{0r}] \end{aligned}$$

where

$$\theta^* = \frac{\theta V_{nm}^{0r}}{\theta V_{nm}^{0r} + (1 - \theta) V_m^{0r}}.$$

The log-linearization is around the point when the agent holds no stocks and is exposed to no home price risk. An analogous formula holds for the right hand side, with the same constants. Using the approximation $\log E \exp(z) \approx Ez + \sigma_z^2/2$ which is exact when z is a normal random variable, we then obtain

$$Er - r_f + \frac{\sigma^2}{2} \approx \theta^* \cdot \text{cov}[r, -v'_{nm}] + (1 - \theta^*) \cdot \text{cov}[r, -v'_m]$$

as in the text. To compute θ^* , denote by W_1^0 wealth in $t = 1$ assuming a safe portfolio and that the home price equals $P_1^0 = \exp \mu_p$. Substituting into (6) and (7) yields

$$\begin{aligned} \frac{V_{nm}^{0r}}{(W_1^0)^{-\gamma}} &= \frac{1 - \mu}{(W_1^0)^{-\gamma}} H_0^{\mu(1-\gamma)} (W_1^0 - P_1^0 H_0)^{-\mu-\gamma+\mu\gamma} = \left(\frac{H_0}{W_1^0}\right)^{\mu(1-\gamma)} (1 - \mu) \left(1 - \frac{P_1^0 H_0}{W_1^0}\right)^{-\mu-\gamma+\mu\gamma} \\ \frac{V_m^{0r}}{(W_1^0)^{-\gamma}} &= \left[\mu^\mu (1 - \mu)^{1-\mu}\right]^{1-\gamma} (P_1^0)^{-\mu(1-\gamma)} \end{aligned}$$

which imply, after some calculations, equation (3) in the text.

Now note that $V'_{nm}(W_1)$ is proportional to $(L_0(1 + R_p) + Y_1 - M_1)^{-\mu-\gamma+\gamma\mu}$. Let $L_0(1 + R_p) + Y_1 - M_1 = L_1$, which we can loglinearize as

$$l_1 \approx k + \eta_1(l + r_p) + \eta_2(y + r_f) + (1 - \eta_1 - \eta_2)(m + r_f)$$

where k is a constant, lowercase letters denote the logs of L , Y and M defined in the text, and

$$\eta_1 = \frac{L}{L + Y - M} \text{ and } \eta_2 = \frac{Y}{L + Y - M}.$$

$V'_m(W_1)$ is proportional to $W_1^{-\gamma} P_1^{(\gamma-1)\mu}$. We can loglinearize $W_1 = L_0(1 + R_p) + P_1 H_0 + Y_1 - M_1$ as

$$w_1 \approx k' + (1 - \rho)\eta_1(l + r_p) + (1 - \rho)\eta_2(y + r_f) + (1 - \rho)(1 - \eta_1 - \eta_2)(m + r_f) + \rho p_1$$

where k' is a different constant and $\rho = PH/(L + PH + Y - M)$ is the housing share in wealth W .

Substituting these expressions into $V'_{nm}(W_1)$ and $V'_m(W_1)$ and then in the Euler equation yields

$$Er - r_f + \frac{\sigma^2}{2} \approx (1 - \theta^*) \gamma \cdot [(1 - \rho) \eta_1 \cdot \alpha \sigma^2 + (\rho - (1 - 1/\gamma) \mu) \text{cov}[r, p_1]] + \theta^* \gamma^c \eta_1 \alpha \sigma^2$$

and hence

$$\alpha \approx \frac{Er - r_f + \sigma^2/2 + (1 - \theta^*) (\mu (\gamma - 1) - \gamma \rho) \text{cov}[r, p_1]}{\theta^* \gamma^c \cdot \sigma^2 \eta_1 + (1 - \theta^*) \gamma \cdot \sigma^2 \eta_1 (1 - \rho)}$$

which gives (3) as desired.

Numerical solution. We use the same numerical techniques as Cocco (2005) to solve the model. The idea is use backward induction and compute continuation values over grids. We approximate the state and choice variables using equal-spaced grids, and the probability density functions of shocks with three-point Gaussian quadratures. In the static specifications, both with random moves and fixed costs, and with and without labor income risk, we compute realized utility over each gridpoint in the state space and then expected utility using numerical integration for each choice of α . To compute utilities for points which do not lie on the grid, we use cubic spline interpolation. In the dynamic model, we perform the same exercise in the last period for each choice of consumption and portfolio, and use the resulting optimal continuation values to solve for the optimal portfolio in the initial period. Using a seven-point Gaussian quadrature does not alter the results.

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.

Cocco, J. F., Gomes F. J. and Maenhout P. J. (2005) “Consumption and Portfolio Choice over the Life Cycle,” *Review of Financial Studies* 18, 491-533.

Davidoff, Thomas (2010). “Home Equity Commitment and Long-Term Care Insurance Demand,” *Journal of Public Economics* 94(1), 44-49.

Davis, J. S., F. Kubler and P. Willen (2006) “Borrowing Costs and the Demand for Equity over the Life Cycle,” *Review of Economics and Statistics* 88, 348-362.

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.

Himmelberg, C., Mayer, C., and Sinai, T. (2005). “Assessing High House Prices: Bubbles, Fundamentals, and Misperceptions,” *Journal of Economic Perspectives* 19(4): 67-92.

Merton, R. C. (1969). “Lifetime Portfolio Selection Under Uncertainty: The Continuous Time Case,” *Review of Economics and Statistics* 51, 247-257.

Mian, A. and Sufi, A. (2011). “House Prices, Home Equity Based Borrowing, and the U.S. Household Leverage Crisis,” *American Economic Review* 101, 2132-2156.

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.

Saiz, A. (2010). “The Geographic Determinants of Housing Supply,” *Quarterly Journal of Economics*, 125(3): 1253-1296.

Shore, S. and Sinai, T. (2010). “Commitment, Risk, and Consumption: Do Birds of a Feather Have Bigger Nests?” *Review of Economics and Statistics*, 92(2), 408–424.

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

Smith, L., Rosen, K., and Fallis, G. (1988). “Recent Developments in Economic Models of Housing Markets,” *Journal of Economic Literature*, 26(1): 29-64.

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.

Yao, R. and Zhang, H. (2005). “Optimal Consumption and Portfolio Choices with Risky Housing and Borrowing Constraints,” *Review of Financial Studies*, 18(1), 197-239.

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