

**Finance and Economics Discussion Series
Divisions of Research & Statistics and Monetary Affairs
Federal Reserve Board, Washington, D.C.**

Housing, Consumption, and Credit Constraints

Andreas Lehnert

2004-63

NOTE: Staff working papers in the Finance and Economics Discussion Series (FEDS) are preliminary materials circulated to stimulate discussion and critical comment. The analysis and conclusions set forth are those of the authors and do not indicate concurrence by other members of the research staff or the Board of Governors. References in publications to the Finance and Economics Discussion Series (other than acknowledgement) should be cleared with the author(s) to protect the tentative character of these papers.

Housing, Consumption, and Credit Constraints^{*}

Andreas Lehnert

Board of Governors of the Federal Reserve System
Washington, DC 20551
(202) 452-3325
Andreas.Lehnert@FRB.GOV

First Version: Sep. 2002
Last Revised: September 29, 2004

^{*}I thank Joshua Gallin, Darrel Cohen, and seminar participants at several institutions for valuable suggestions. Mary F. DiCarlantonio for excellent research assistance. Any remaining errors are my own. The opinions, analysis and conclusions in this paper are my own and do not necessarily represent those of the Board of Governors of the Federal Reserve System or its staff.

Housing, Consumption, and Credit Constraints

Abstract

I test the credit-market effects of housing wealth shocks by estimating the consumption elasticity of house price shocks among households in different age quintiles. Younger households face faster expected income growth and hence would like to borrow more than older households. I estimate consumption elasticities from housing wealth by age quintile to be $\{4, 0, 3, 8, 3\}$ percent. As predicted by theory, the youngest group has a higher elasticity of consumption than the next two age quintiles. That the consumption of the age quintile on the verge of retirement is responsive to housing wealth is also not surprising: I show that these households are likeliest to “downsize” their house and thus realize any capital gains.

Journal of Economic Literature classification numbers: D12, D91, E21, G21

Keywords: Consumption, wealth effect, housing, house prices

1 Introduction

In the U.S., housing is an important part of the typical American's balance sheet. Moreover, house values have continued to climb even as the stock market has stagnated. Thus, it is particularly important to understand the effect of housing wealth on consumption.

Housing wealth can affect consumption through many channels. For example, homeowners can move to cheaper quarters and realize some of the equity from their previous dwelling or they can assume added debt backed by the wealth of their house. Moreover, households may view housing wealth more as a buffer stock of wealth to be used in an emergency or to finance a specific expenditure. Typically, however, most households are not experiencing these sorts of contingencies, that is, they are not moving into cheaper houses, experiencing economic stress, or paying a child's college tuition. How then do households react to quotidian increases or decreases in their housing wealth; that is, to housing wealth shocks that arrive in the normal course of their lives?

Further, even if such quotidian increases in housing wealth do affect consumption, it is far from clear that this can translate into increased aggregate consumption. As a matter of basic accounting, increases in house prices are exactly offset by an increase in the user cost of housing. From this accounting view, house prices can rise indefinitely without affecting the aggregate consumption of non-housing goods and services. Indeed, in a simple model with a fixed supply of housing and increasing labor productivity, house prices must rise indefinitely. Without some

kind of frictions, in such a model housing wealth shocks *per se* will be irrelevant to aggregate macroeconomic conditions.¹ Yet to claim that housing wealth does not matter at all for aggregate consumption flies in the face of intuition, not to mention received wisdom.

Housing wealth could affect aggregate consumption in its role as collateral; that is, agents who own (and have equity in) their homes can more easily pledge to repay loans, hence overcoming commitment problems in credit markets. This idea is exploited by Lustig and Van Nieuwerburgh (2002), who study the effect of rising house values on stock market returns. In this view, increases in housing wealth can cause decreases in the aggregate savings rate because of the increased collateral value of housing, and its associated relaxation of borrowing constraints.

This “credit channel” of housing wealth can have a further effect if one believes that financial developments over the past decade have made borrowing against home equity easier and cheaper. Thus, as a result of financial innovation, we could be in the transition from an old equilibrium, in which housing wealth did not matter much, to a new equilibrium in which housing wealth directly relieves borrowing constraints. This hypothesis, however, awaits formal econometric evidence.

If housing wealth relaxes borrowing constraints, we would expect its effects to appear first in households’ precautionary saving. Carroll, Dynan, and Krane (1999) find that, among households with higher labor income risk, precautionary

¹See Lehnert and Pence (2003) for one such model that emphasizes the role for foreclosure law on housing finance markets and house prices.

savings was higher, but that most of this increase was in the form of housing wealth. One explanation for this could be that housing wealth has become easier to tap in the event of a family emergency.²

Another, though related, explanation is that housing wealth enjoys special tax status and bankruptcy court status. The treatment of home equity in bankruptcy proceedings is relevant to households considering their optimal plan following severe shocks, such as extended illness or job loss.³

Hurst and Stafford (2002) use the PSID to show that mortgage *refinancing* plays a major role in consumption insurance. As refinancing costs fall, the buffer-stock role of home equity would presumably be greater. That is, homeowners today may be more confident in their ability to borrow against their housing wealth (e.g. by refinancing) than in the past. As a result their savings in other, more liquid, assets may decline. The volume of refinancing, in particular the quantity of so-called “cash out” refinancing, is a measure of how actively households are re-allocating their portfolios or tapping their accumulated home equity for a specific purpose.⁴ In a refinancing, households are pursuing an active means of realizing their home equity gains. However, households may also react passively to home

²For more on the large buffer-stock savings literature, see Carroll (2000, 1994, 1992, 1996, and 1997a,b), Carroll and Samwick (1998), Hubbard, Skinner, and Zeldes (1994), and Engen (1993).

³Lehnert and Maki (2002) study the effect of bankruptcy law on household consumption; they find that—especially among homeowners—bankruptcy law affects the consumption response to income shocks.

⁴There is no generally agreed-upon definition of “cash out” refinancing. However, most industry experts agree that it involves an increase in the outstanding mortgage principal by more than 2.5 percent. Informally, the homeowner walks away from the refinancing with a check, which he or she then puts into another form of asset (portfolio reallocation) or uses to purchase goods or services (specific consumption purpose).

equity gains. This paper is an attempt to pin down the nature and importance of this passive channel.

The relationship of house prices to consumption has been an active area of research interest for more than a decade. Several studies have addressed the effect of house price increases on the consumption of young renters who are, one presumes, saving for house downpayments. Among these are Sheiner (1995) and Engelhardt (1994). Under one view, renters should increase saving in the face of house price increases in order to better afford a down payment; under the other view, renters give up trying to buy a house and succumb to – in the memorable phrase – “the consumption of despair.”

This paper follows the seminal study by Skinner (1996), who also uses the PSID to test for housing wealth effects on consumption. He finds that wealth effects are greatest among young households, and that these younger households decrease their active savings in response to housing wealth shocks.

These studies addressed the effects of housing wealth shocks on individual households. Fewer studies have tackled the problem of estimating the effects of housing wealth shocks on aggregate consumption. Studies using data on aggregate U.S. consumption tend to have a hard time separating the effect of housing wealth shocks from stock price and other wealth shocks. Case, Quigley and Shiller (2001) use the extra variation provided by U.S. states, as well as a cross-country panel, to identify the separate effects of stock and housing wealth. The authors find that the marginal propensity to consume (MPC) out of housing wealth far exceeds the

MPC out of stock-market wealth.⁵

In this paper, I test whether the consumption impact of house price shocks is greater among credit constrained households than among other households. I identify credit constrained households by *age*. Young households face steeper income increases than older households, and hence a greater incentive to borrow. Conversely, they face a decreased desire for saving.

Specifically, I use household-level data from the 1968-1993 waves of the *Panel Study of Income Dynamics* (PSID) to address this question.⁶ In this study, I exclude households that experienced severe distress (such as divorce or death of the head) in the previous year; that changed composition (such as adding or losing a child at home); and, perhaps most importantly, that moved. These are precisely those households that might be expected to consume a large portion of any accumulated housing wealth. Instead of studying these households, I study a sample of households that are stable over (at least) the current and previous year.

Among this sample of stable homeowners I find an average consumption elasticity out of housing wealth gains of between 0.04 and 0.05. These elasticities translate into marginal propensities to consume of between 2 and 3 cents per dollar. When I split my sample into groups by age I find striking differences in these estimated effects. The effect of housing wealth on consumption is greatest for

⁵Case, Quigley, and Shiller's finding that housing wealth affects consumption more than stock-market wealth may be prey to Poterba's (1991) observation that house prices capitalize expected future gains in labor income. Thus, households in a region receive good economic news and *simultaneously* increase their non-durable consumption while bidding up local house prices.

⁶These waves constitute the current final release of the PSID; data from 1994–1999 are available only in “early release” form. Early release data do not contain accuracy codes or sample weights.

households in late middle age (ages 52 through 62); however, the next-most affected age group are the *youngest* households. Other age groups show smaller effects.

The fact that younger households react so strongly to housing wealth gains might appear, at first, to be something of a puzzle. After all, young homeowners are most likely to move into a bigger house as their family size increases. House price appreciation is not necessarily good news for this group.

There are two reasons for younger households to consume a greater share of their housing wealth gains. First, young households are more likely to move than older households; as a result, they are more likely to realize their housing wealth gains.

Second, young households may be acting as buffer-stock or liquidity constrained consumers. Younger households face faster-growing permanent incomes than do older households. As a result they would like to borrow against their future income gains, or, failing that, at least not save very much. Thus, they are more likely to consume shocks to any form of wealth, including housing wealth. As evidence for this view, Gourinchas and Parker (2002) show that households generally act like buffer stock consumers until around age 40, when they begin to act more like traditional life-cycle consumers.

Finally, it's important to note that the estimated elasticity among younger households declines if I do not throw households that move or change family composition out of my dataset. This is probably because changing family composition and moving affect consumption in ways that cannot easily be controlled

for.

The fact that households in late middle age (ages 52–62) are most sensitive to housing wealth gains is no surprise. These households are on the eve of retirement and making their post-retirement housing choices. Evidence from the SCF and CPS shows that beyond this age, households generally live in less valuable houses. In addition, they ought to have largely completed saving for retirement, so that additions to housing wealth translate into expected realized wealth gains.

Researchers agree that households are likely to spend large portions of *realized* home equity gains. However, little was known about households' propensity to consume out of passive housing wealth gains (as opposed to active housing wealth gains, which are realized in sales or in a refinancing). In this paper I find reasonable estimates of the marginal propensity to consume out of housing wealth gains. Further, I find a strongly non-monotone age pattern to these estimates. However, the age pattern also fits with economic theory.

The rest of this paper is organized as follows: In section 2 I briefly review the aggregate importance of housing wealth; in section 3 I present evidence from the PSID and the CPS about the effect of age on housing demand; in section 4 I present estimated consumption elasticities of housing wealth; in section 5 I translate these elasticities into marginal propensities to consume using weights from the Consumer Expenditure Survey (CE); finally, I briefly conclude and describe ongoing refinements to this research in section 6.

I include several appendices to this paper. In appendix A I review the PSID and present precisely the steps I used to arrive at my sample of “stable households.”

In appendix B I present sample statistics from my dataset. Finally, in appendix C I present exhaustive results from alternative specifications and different samples from the PSID. Almost all of my results are robust to all specification changes and sample selections. However, my finding that younger households have high consumption elasticities can be weakened if one does not eliminate observations from the sample in which households move or experience compositional changes.

2 Aggregate Trends in Housing Wealth

In the past decade housing wealth has grown as a share of household portfolios; moreover, these gains have occurred against the backdrop of a historic rise in the rate of home ownership. Beginning in late 2000 and continuing through the present, residential real estate has appreciated faster than stock prices. Indeed, gains in housing wealth have been credited with propping up consumption even as the stock market has slumped.

Figure 1 shows the relative performance of residential real estate and the stock market. Although over the longer horizon, equities have heavily outperformed residential real estate, over the period 2000–2004 house prices have risen at a historically rapid pace while equity prices have slumped. The only comparable period is the mid- to late-1970s, when house prices outperformed stock prices for about five years.

Figure 2 shows the relative importance of equities and residential real estate in the household sector's aggregate balance sheet. Real estate has traditionally made

up the bulk of the household sector's assets; the large share accounted for by equities in the second half of 1990s was something of a historical anomaly. In the past year, as the prices of the two asset classes moved in opposite directions, housing has once again become the dominant asset on the household balance sheet.

3 Housing and Income Growth by Age

House size and income growth are closely related to household age. Younger households are likelier to anticipate fast income growth and are also likelier to anticipate increased demand for housing services. In this section I present evidence on the empirical relationship between housing consumption and age and between income growth and age.

3.1 Housing Consumption and Age

Housing choices are closely tied to household formation and size; younger, childless, couples will demand a different mix of housing services than older couples with children. One might imagine that younger households would be more likely to increase their housing stock (to “upsized”) and older couples more likely to decrease their housing stock (to “downsize”); however, this demand-driven trend is somewhat obscured by the fact that younger households are much more likely to move than older households. Sheiner and Weil (1992) conducted the first study of the housing wealth of the old; they found that older households were more likely to downsize after a severe shock, such as sudden medical expenses or death of a

spouse. They found that this downsizing occurred relatively gradually. I also find that older households are quite likely to remain owner-occupiers; however, I do find that older households are much more likely to be downsizers than younger households, conditional on moving.

I used all waves of the PSID (1968–1999) to construct a dataset of the tenure status of household heads. The PSID attempts to follow households over time, so it was possible for most households to determine whether, once a household moved, it increased or decreased the value of its house. Also, I classed as upsizers (downsizers) those households that moved into rented quarters (owned quarters) after having previously owned (rented) their home. I then determined the probability of moving by age and the probability of upsizing (downsizing) by age conditional on having moved.

As a measure of the demand for housing services, figure 3 presents the average and median number of rooms for each household by the household head's age. Dwelling size increases with age until age 45, when it starts to decline.

As shown in the top panel of figure 4, households in lowest age quintile (between 25 and 34) have roughly twice the probability of moving per year as households in older quintiles. Conditional on moving, however, as shown by the lower panel of figure 4, younger households are two to three times more likely to upsize than older households.

My sample of PSID household-heads does not include household heads that moved in with their children or otherwise ceased to head their household. If an elderly parent moves in with his son, the son is considered the head and the parent

drops out of my sample. Thus I might understate the true decrease in housing demand by age. I turned to the 1976–2001 waves of the March *Current Population Survey* to get a representative picture of housing demand by person’s age. Controlling for cohort, year and regional effects, as a person ages the conditional probability of being a household head (or spouse) and a homeowner begins to fall off. The peak occurs around age 65; after this point people are less likely to head their household and to own their homes. However, it is worth noting that at age 90, people were as likely to be homeowner/heads as 45 year-olds.

4 Consumption Elasticities for Homeowners

The PSID collects information about household expenditures on food, house value (if owned) and household income sources. In addition, the PSID typically follows households for several years. In this section I describe how I use the PSID to estimate consumption elasticities out of housing wealth. I find economically and statistically significant values; moreover, I find that these values differ dramatically by age group, with the youngest and the oldest households having the highest elasticities and middle-aged households having the lowest elasticities.

I constructed a sample of “stable households” from the PSID as my base dataset. Appendix A has more information on the data and sample selection procedure, appendix B presents a series of sample statistics for the dataset and appendix C describes results using alternative specifications.

4.1 Specification

Let $C_{i,t}$ denote the real food consumption of household i in year t , $V_{i,t}$ denote the real value of the house and $Z_{i,t}$ denote a vector of additional household-level variables. The specification of interest is:

$$(1) \quad \Delta \log(C_{i,t}) = \alpha \Delta \log(V_{i,t}) + Z_{i,t} \Gamma + u_{i,t}.$$

The estimated parameter α gives the elasticity of food consumption out of housing wealth. Although most of my additional controls are standard (household composition, head's age, year fixed effects, change in family income, etc), three are relatively novel and merit special mention.

First, I include measures of permanent labor income (more precisely, permanent income growth). Because the PSID allows several observations of a single household, I can determine whether a given income shock is transitory or permanent. Household consumption ought to react more to permanent than transitory shocks. Let $\ell_{i,t}$ denote the log real labor income of the household head; for each household i I construct a time series of growth rates:

$$\mathbf{L}_i \equiv \{\ell_{i,t} - \ell_{i,t-1}\}_{t=2}^{T_i}.$$

I then construct the sample average, \bar{L}_i , for each household, as well as indicator variables for growth rates that are in the top and bottom percentile of the sample. Households with faster-growing permanent incomes, here proxied by higher val-

ues of \bar{L}_i , should have faster-growing consumption as well. As Poterba (1991) noted, house prices, like all asset prices, respond quickly to news about economic conditions. Controlling for the household's permanent income growth prevents house price gains due to improving local labor market conditions from spuriously affecting consumption growth.

Second, I included fixed effects for the household's state of residence in year t . These state-level effects remove common elements in consumption growth due to geography; in addition, they control for peculiarities due to conditions in unusual states with few residents. For example, the final dataset contained only three observations from Montana; eight other states were represented by fewer than 50 observations.⁷

Third, I included the state-level growth rate of home prices. These data are only available from 1976; thus including this variable required throwing out several early years of my sample. For this reason I generally report results with and without the extra variable.

I refer to the specification without any of the extra controls as the “base specification.” I then added the three extra sets of control variables in all possible combinations, except that I always included the state fixed effects with state house price growth.

⁷In ascending order: North Dakota (14) Wyoming (19), Idaho (22), Hawaii (26), Rhode Island (33), and New Mexico (36).

4.2 Results

Table 1 reports the estimated elasticity, $\hat{\alpha}$, and the standard error of the estimate for all specifications. The estimated elasticity does not vary with the inclusion of state fixed effects or permanent income controls, remaining constant at 3.94 percent. This estimate is different from zero at the 0.1 percent level of significance (or better). When the sample is constrained to 1976 and forward, and state-level house price growth is included, the estimated elasticity grows to about 4.7 percent.

Because I am interested in how these elasticities vary by age, I divided the sample into quintiles by age and reestimated the models. I show the point estimates and 90 percent confidence regions for the base specification and the specification with all possible controls in figure 7.

Note the striking pattern: The elasticity of consumption for the youngest quintile is higher than for the next two older quintiles. This is consistent with standard economic theories in two respects. First, younger households are more likely to be liquidity constrained and thus use wealth purely as a buffer stock, while older households use wealth for life-cycle reasons. Gourinchas and Parker (2002) find evidence that younger households do indeed act as buffer-stock consumers while older households act more as life-cycle consumers.

A second reason for the higher elasticity of younger consumers is that, as shown in figure 4, younger households are much more likely to move, thus potentially realizing some of their housing wealth gains.

Consistently, though, the highest elasticity of consumption is among the second oldest quintile, the group aged 52 through 62; examining figures 4 and 5 it

is precisely this age group which will begin downsizing and realizing some of its housing wealth gains.

5 Marginal Propensities to Consume out of Housing Wealth

The PSID survey instruments ask about several measures of consumption, including food at home, food away from home, rent, house value and utilities. Although these are important components of total expenditures, they are far from the total. Moreover, there may be systematic differences in the translation from these components to total expenditures by household cohort and demographics. The most popular method of imputing total expenditures was first proposed by Skinner (1987). More recently, Blundell, Pistaferri, and Preston (2002) have proposed inverting estimated Engel curves. Here, I follow the Skinner procedure, although I have also used the Blundell *et al* procedure, with very similar results.

Skinner (1987) suggests using the limited information on consumption in the PSID to infer total expenditures for each household. The Consumer Expenditure Survey (CE) tracks spending on a variety of goods and services, including those also tracked by the PSID. I deviate from the standard procedure by not using information on house value in the PSID to impute total spending. Doing so would introduce a mechanical dependence between house values and consumption. Table 2 gives the regression coefficients from the modified Skinner procedure.

As can be seen, excluding house values degrades the fit substantially. On the

other hand, these data (which are from repeated cross-sections) imply an elasticity of consumption out of housing wealth of between 0.22 and 0.25; eliminating such a mechanical result is worth the price in precision.

One might wonder about the value of including other observed household characteristics, such as the age of the head, number of children and female labor force participation. All of these variables (and others) in fact do enter Skinner-style specifications significantly; however they do not improve the fit of the regression measurably. Further, because I am interested in the response of consumption at different ages, I avoid using household information that is correlated with age (such as female labor participation or the number of children living at home) to construct my measures of consumption.

Using the weights from table 2 I constructed, for each household in each year, imputed levels of non-durable and total consumption. I then re-ran the specifications from section 4 using these imputed levels of consumption. The results (along with other pertinent information) are shown in table 3.

The estimated elasticities based on the imputed consumption measures show the same general pattern as those based on food consumption. The youngest quintile has an elasticity of between 3 and 4 percent, the second-oldest quintile (ages 52-62) has the highest elasticity, of between 3.4 and 4.7 percent and other quintiles have lower elasticities. However, the age pattern of these elasticities is less pronounced than those estimated using food consumption.

The marginal propensity to consume is just the estimated elasticity times the average ratio of consumption to house value. For age quintile j the MPC can be

computed as:

$$MPC_j = \hat{\alpha}_j \frac{1}{N} \sum_{i=1}^{N_j} \frac{C_i}{V_i}.$$

The average ratio of consumption to house value for each quintile is also shown in table 3. Notice that older households consume less relative to the value of their house than do younger households; this is because younger households (as we have discussed) typically live in cheaper houses than do older households. Thus older households will typically spend less of an extra dollar in housing wealth than younger households.

The estimated MPCs shown in table 3 reflect these two age-dependent effects. The youngest households have the highest elasticities and the ratios of consumption to house value; thus they have the highest MPC. Households in the next age quintile (age 35–42) have zero elasticities of consumption with respect to housing wealth, so their MPCs are also essentially zero. The last three age quintiles all have about the same MPC out of housing wealth, somewhere between 2 and 3.9 percent.

6 Conclusion

In this paper I used the PSID to estimate the effect of passive housing wealth gains on consumption. I explicitly excluded those observations in which the household moved, as well as other observations in which consumption growth would react

to household-level changes. Among this group of stable households, I estimated a total sample MPC out of housing wealth of between 1.9 and 3.1 cents per dollar. This value, though, concealed variations across age quintiles. This variation could be explained by reasonable economic forces: younger households are both liquidity constrained and more likely to move than older households. Households with the highest sensitivity to housing wealth gains (those aged 52–62) are precisely those preparing to retire and thus likeliest to move into smaller houses.

References

- Blundell, R., L. Pistaferri, and I. Preston (2002). Partial insurance, information, and consumption dynamics. Manuscript, Department of Economics, Stanford University, Palo Alto CA.
- Carroll, C. D. (1992). Buffer stock saving: Some macroeconomic evidence. *Brookings Papers on Economic Activity* 1992(2), 61–156.
- Carroll, C. D. (1994, February). How does future income affect current consumption? *Quarterly Journal of Economics* 109(1), 111–47.
- Carroll, C. D. (1996). Buffer stock saving: Some theory. Manuscript, Department of Economics, The John Hopkins University.
- Carroll, C. D. (1997a). Buffer-stock saving and the life-cycle/permanent income hypothesis. *Quarterly Journal of Economics* 112(1), 1–56.
- Carroll, C. D. (1997b). Unemployment expectations, jumping (S, s) triggers, and household balance sheets. In B. S. Bernanke and J. Rotemberg (Eds.), *NBER Macroeconomics Annual, 1997*. Cambridge, MA: MIT Press.
- Carroll, C. D. (2000). Requiem for the representative consumer? Aggregate implications of microeconomic consumption behavior. *American Economic Review* 90(2), 110–15.
- Carroll, C. D., K. E. Dynan, and S. D. Krane (1999). Unemployment risk and precautionary wealth: Evidence from households' balance sheets. Finance

- and Economics Discussion Series 1999-15, Federal Reserve Board.
- Carroll, C. D. and A. A. Samwick (1998). How important is precautionary saving? *Review of Economics and Statistics* 80(3), 410–19.
- Case, K. E., R. J. Shiller, and J. M. Quigley (2001, November). Comparing wealth effects: The stock market versus the housing market. NBER Working Paper 8606, National Bureau of Economic Research.
- Engelhardt, G. V. (1994, September). House prices and the decision to save for down payments. *Journal of Urban Economics* 36(2), 209–37.
- Engen, E. (1993). Consumption and saving in a life-cycle model with stochastic earnings and uncertain lifespan. Manuscript, Federal Reserve Board, Washington D.C.
- Gourinchas, P.-O. and J. A. Parker (2002). Consumption over the life cycle. *Econometrica* 70(1), 47–89.
- Hubbard, R. G., J. Skinner, and S. P. Zeldes (1994, May). Expanding the life-cycle model: Precautionary saving and public policy. *American Economic Review* 84(2), 174–79.
- Hurst, E. and F. Stafford (2002, August). Home is where the equity is: Liquidity constraints, refinancing and consumption. Graduate School of Business, University of Chicago, Chicago IL.
- Lehnert, A. and D. M. Maki (2002). Consumption, debt, and portfolio choice: Testing the effects of bankruptcy law. Finance and Economics Discussion Series 2002-14, Federal Reserve Board, Washington, D.C.
- Lehnert, A. and K. M. Pence (2003). The price of protection: Foreclosure law and house prices. Unpublished manuscript, Federal Reserve Board, Washington D.C.
- Lustig, H. and S. Van Nieuwerburgh (2002, November). Housing collateral, consumption insurance and risk premia. Manuscript, University of Chicago, Department of Economics, Chicago IL.
- Poterba, J. M. (1991). House price dynamics: The role of tax policy and demography. *Brookings Papers on Economic Activity* 1991(2), 143–83.
- Sheiner, L. (1995, July). Housing prices and the savings of renters. *Journal of Urban Economics* 38(1), 94–125.
- Sheiner, L. and D. N. Weil (1992, July). The housing wealth of the aged. NBER Working Paper 4115, National Bureau of Economic Research.

- Skinner, J. S. (1987). A superior measure of consumption from the Panel Study of Income Dynamics. *Economics Letters* 23, 213–16.
- Skinner, J. S. (1996). Is housing wealth a sideshow? In *Advances in the Economics of Aging*, National Bureau of Economic Research Report, pp. 241–268. Chicago, IL: University of Chicago Press.

TABLE 1: Estimated Consumption Elasticities

Specification	Estimate	Standard Error
<i>Without Permanent Income Controls</i>		
Base	0.0394	0.0105
State Dummies	0.0390	0.0105
State Dummies+State HPI Growth	0.0471	0.0118
<i>With Permanent Income Controls</i>		
Base	0.0394	0.0105
State Dummies	0.0390	0.0105
State Dummies+State HPI Growth	0.0472	0.0118

NOTE. Table gives OLS estimates of the elasticity of consumption, $\hat{\alpha}$, with respect to real house value gains among a sample of stable homeowners. Without state house price controls there were 19,316 observations; with the state house price controls there were 15,988 observations.

TABLE 2: Consumption Weights from the Consumer Expenditure Survey (CE)

Regressors	Consumption Measure			
	Total		Non-Durable	
Log food away from home .	0.2156 (0.0031)	0.2746 (0.0030)	0.1849 (0.0023)	0.2466 (0.0025)
Log food at home	0.2924 (0.0063)	0.3612 (0.0065)	0.3101 (0.0047)	0.3921 (0.0055)
Log rent equivalent	0.2246 (0.0034)		0.2504 (0.0025)	
Constant	4.4921 (0.0522)	5.4294 (0.0528)	4.1097 (0.0039)	5.1463 (0.0450)
R^2	0.6179	0.5021	0.7420	0.5608

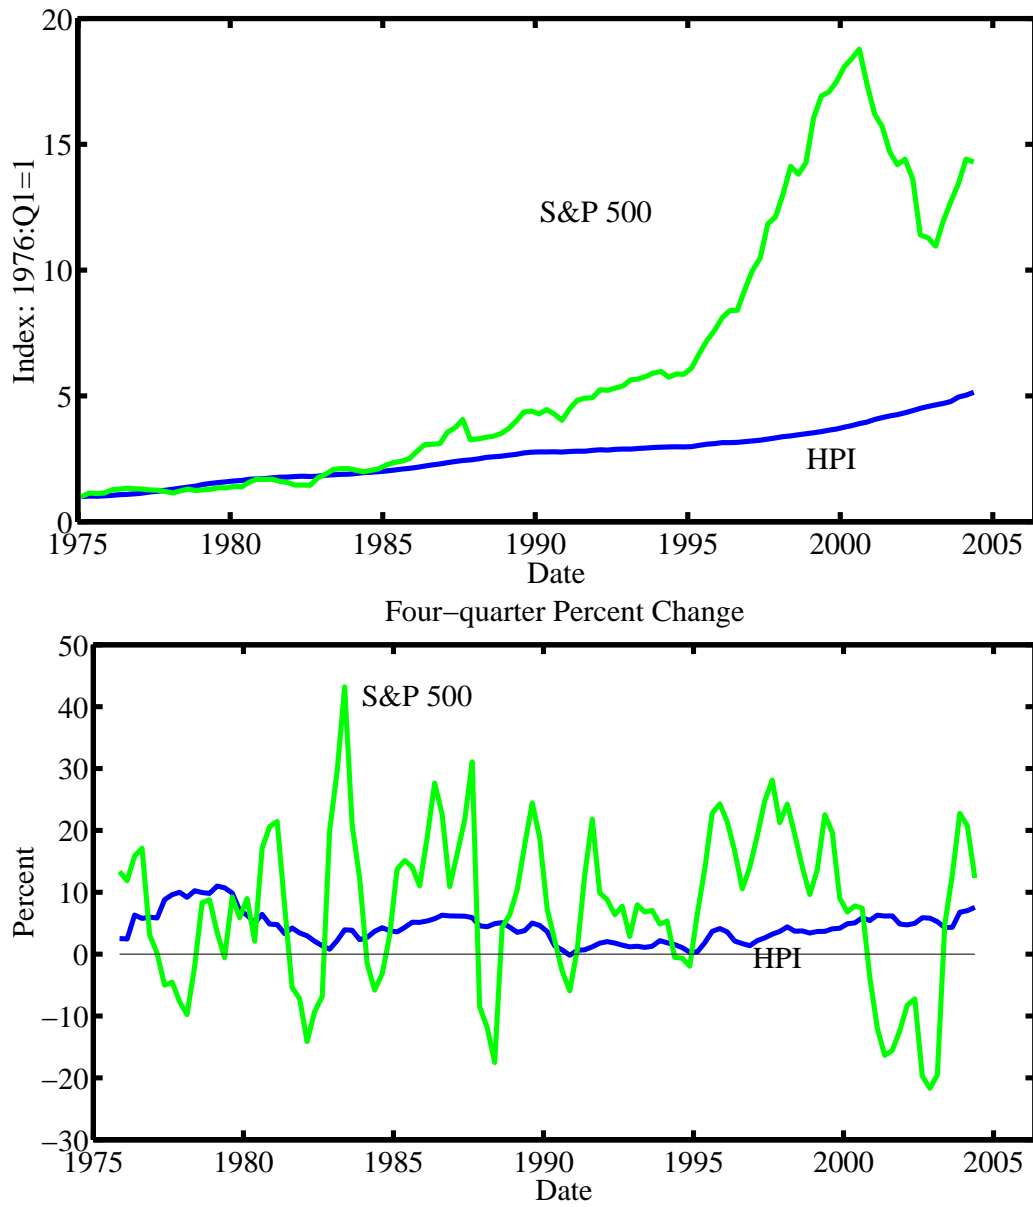
NOTE. Table gives coefficients from an OLS regression of the log of the indicated consumption measure on the indicated regressors; all values are real, deflated by the appropriate deflators. Standard errors are in parentheses. Data are from the Consumer Expenditure Survey, 1984–2001.

TABLE 3: Estimated MPCs out of Housing Wealth

	Age Quintiles					
	All	25–34	35–42	43–51	52–62	63–95
Total Consumption						
<i>Base</i>						
Elasticity (percent)	2.64	3.05	0.00	2.91	3.36	2.90
C/V	0.91	1.04	0.92	0.88	0.83	0.87
MPC (percent)	2.41	3.18	0.00	2.57	2.80	2.52
<i>Permanent Income Controls + State Dummies + State House Price Gains</i>						
Elasticity (percent)	3.43	3.75	0.01	3.93	4.67	2.97
C/V	0.91	1.04	0.92	0.88	0.83	0.87
MPC (percent)	3.13	3.91	0.01	3.47	3.89	2.57
Non-Durable Consumption						
<i>Base</i>						
Elasticity (percent)	2.65	3.06	0.00	2.92	3.37	2.92
C/V	0.70	0.80	0.71	0.68	0.64	0.67
MPC (percent)	1.87	2.46	0.00	1.99	2.17	1.95
<i>Permanent Income Controls + State Dummies + State House Price Gains</i>						
Elasticity (percent)	3.44	3.77	0.01	3.95	4.69	2.98
C/V	0.70	0.80	0.71	0.68	0.64	0.67
MPC (percent)	2.41	3.02	0.00	2.68	3.00	1.99


NOTE. Table gives data necessary to compute MPCs out of housing wealth from the PSID. The top panel gives results using imputed *total* consumption and the bottom panel gives results using imputed *non-durable* consumption. Rows give the estimated elasticities, the average ratio of (imputed) consumption to house value and their product (the MPC).

FIGURE 1: Equity and Residential Real Estate Prices in the U.S. 1975–2004




SOURCE. Office of Federal Housing Enterprise Oversight (OFHEO) and Standard & Poors.

FIGURE 2: Aggregate U.S. Household Portfolios 1960–2004



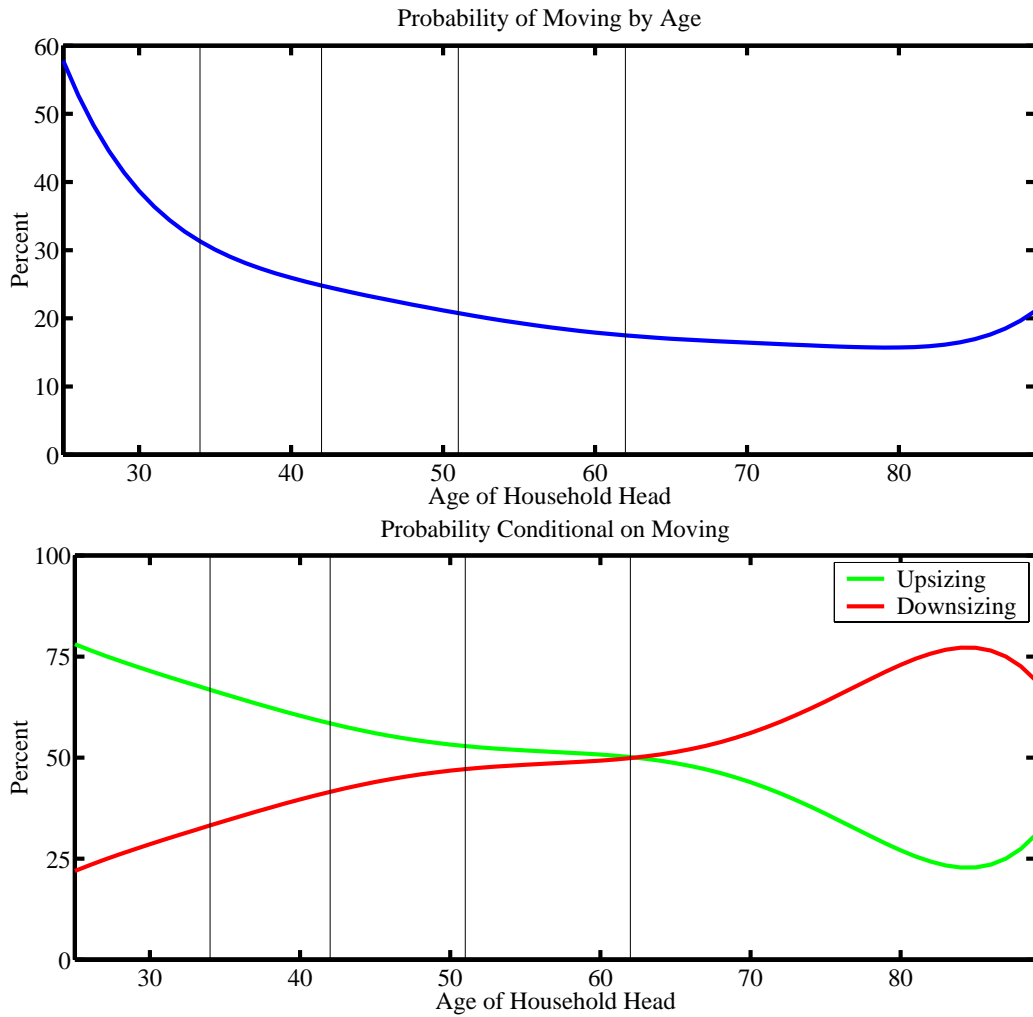
SOURCE. Flow of funds accounts of the U.S.

FIGURE 3: Demand For Housing Services by Age




SOURCE. Panel Study of Income Dynamics, 1968–1999.

FIGURE 4: Homeownership Decisions By Head's Age



NOTE. Figures give the probability of moving (top panel) and the probability of upsizing or downsizing conditional on moving (bottom panel). Data are from the PSID 1968–1999; graphs show linear probability models fit against sixth-order Chebyshev polynomials in the head's age. Regressions also included a full set of year, cohort and state fixed-effects.

FIGURE 5: Head/Homeownership Rates By Age



NOTE. Figure gives the probability of being a household head (or spouse) and a home owner conditional on age; data are from the March CPS, 1976–2001. Plots show the results of probit regressions on a full set of cohort, year and regional fixed effects and a sixth-order Chebyshev polynomial in the head’s age.

FIGURE 6: Income by Age

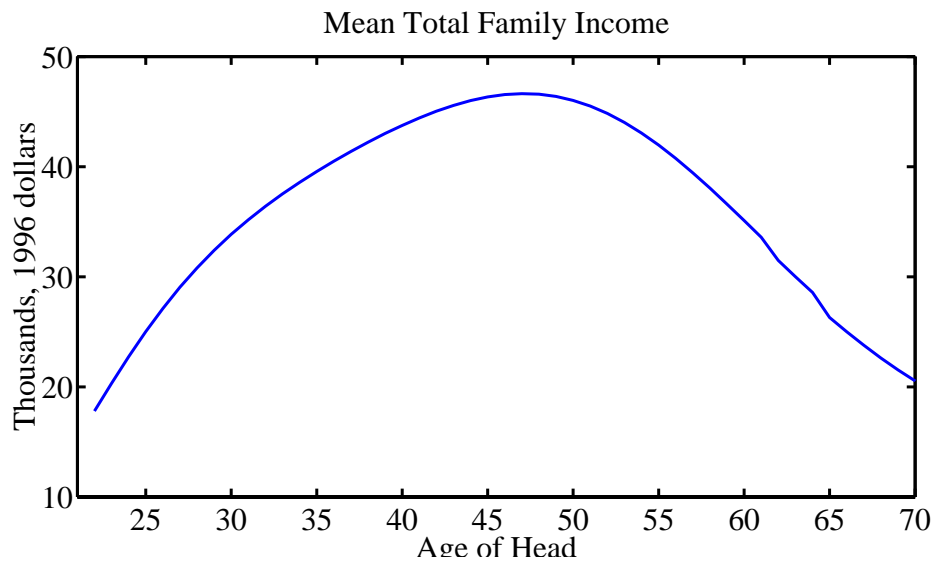



FIGURE 7: Elasticities by Age Quintile



NOTE. Figures give the consumption elasticities estimated via OLS for each age quintile. The top panel gives results from the base specification, the bottom panel from the specification containing all possible controls. The lighter lines give the 90 percent confidence interval.

A Data Construction

I constructed the sample for use in my analysis in two steps. In the first step, I picked a sample of “stable households” out of the PSID. For these households I have information on food consumption (except in certain years, see below), labor and non-labor income, and house value or rental payments. In the second step, I picked which years for each stable household I could use in my regressions. There were a few reasons to throw out particular years for a given household; if the household had moved the previous year, if its size had changed, or, most importantly, if consumption information had not been collected in the previous year.

A.1 Overview of Methodology

I began with both the final release (1968–1993 waves) and the current early release of the PSID (1994–1999 waves). I use information from the early release only to correct missing information from the final release. The PSID began as a representative sample of households in 1968; it has attempted to follow all split-offs (stemming from new household formation by now-adult children of the original sample, divorces among the original sample and so on). I follow the standard practice in restricting the sample to household heads and their families only. Heads need not be male; they are the primary earner. I also eliminate the special poverty subsample (this is also standard practice).

The results of this initial draw is a sample of 113,659 responses from 10,412 families that completed the PSID annual survey anywhere from one to 31 times (the PSID was not administered in 1998). Families can enter and exit the PSID in any year; the most common pattern (705 families) is to answer each PSID survey instrument (i.e. to be in the sample for the full 31 possible years). This is the basic sample that I used to construct my final dataset.

My sample selection worked in this order:

1. I dropped a number of household-years for reasons outlined below; however, if a given household violated the criteria in a given year, other observations from the household were kept.
2. I then required that *households* remain in my sample for at least six consecutive years. Each household could remain in my sample only once. I call this my sample of *stable households*.

3. Of all possible remaining household-years, I determined which ones were associated with valid growth rates in consumption and house value.

Notice that I used two types of criteria: The first type (used in step one) excluded a particular household-year from consideration. This could then in turn prevent the household from entering the final sample (if for example it broke up the household's sample into blocks of five or fewer observations). The second set of criteria (used in step 3) could eliminate any number of observations from a stable household.

A.2 Initial Selection

The initial sample selection eliminated household-years that would in general render the observation unusable. Note that although I am ultimately interested only in homeowners, I can use information from the same household in years when it was a renter. For example, certain variables are not collected each year, but are unlikely to change over time (such as the respondent's race).

Table A.1 gives the steps in my initial sample selection; the table shows how many observations (defined as household year combinations) were dropped at each step and how many separate households remained. I dropped households who neither owned nor rented their dwelling place (mainly students), those with household heads aged 24 or below, those who were not married, those with total family income before taxes of zero (this is an invalid code), and those who lived abroad. In the end, I required households to have completed at least six consecutive valid interviews in the period covered by the final release, 1968–1993. Variables of interest from the early release have not been checked for accuracy, and the PSID warns against their use.

A.3 Selecting Household-Year Combinations

Basic Restrictions

The selection procedure described in table A.1 left me with an unbalanced panel of 39,642 potential household-year observations on 2,832 stable households. Not all of these household-year combinations can be used in the consumption regressions. The PSID did not collect consumption information in its 1968, 1973, 1988 and 1989 waves, nor did it collect information on utilities paid in 1973–1976, 1982, 1988–1993.

The PSID collects retrospective information on income, hours worked and employment status, so that answers included in the 1990 survey responses refer to the household’s experience in 1989. In constructing my final dataset, I shifted all of these retrospective variables back one calendar year. The food consumption questions may or may not be retrospective in nature. Most researchers who use the PSID have concluded that they refer to average food expenditure in the month or quarter in which the survey was given. I follow this practice.

Let $v_{i,t}$ denote the real value of household i ’s house as reported in the interview year t and $c_{i,t}$ denote the real value of food consumption by i in year t . Here “food consumption” is defined as the annualized food expenditures on food at home and food away from home, plus the net benefit from food stamps, if any. Nominal food expenditures are deflated by the relevant deflator at the interview month. The primary relationship of interest is between $\Delta c_{i,t}$ and $\Delta v_{i,t}$, defined as:

$$\Delta c_{i,t} = \log(c_{i,t}) - \log(c_{i,t-1}),$$

$$\Delta v_{i,t} = \log(v_{i,t}) - \log(v_{i,t-1}).$$

Thus we must eliminate not only those years in which the PSID does not collect information on consumption, we have to eliminate the next year’s observations as well.

Additional Restrictions

The restrictions described above are the minimum necessary to produce a dataset in which every observation has a value for consumption growth and for house value growth. I propose an additional set of restrictions designed (1) to isolate the effect of house value changes among households that do not move; (2) to eliminate households that might change consumption for reasons unrelated to house appreciation, such as changes in family size; and (3) to improve the fit of a linear model by eliminating outliers.

In my remaining sample, 1,246 observations were homeowners who moved to another owned home in the previous year. Because the present study is designed to *ignore* the role of seller’s equity extraction, it seems reasonable to remove these observations. A further 3,970 household-years represented households that changed composition in some important way. The reported change in consumption of these families will likely reflect in large part these changes; further, controlling for such changes involves essentially taking out the conditional

mean change in consumption growth for households that experienced the same change. With so few changes in each category, these conditional means are likely to be badly estimated.

Finally, my specifications involve a linear relationship between consumption growth and house value growth. The top and bottom percentile of the sample represent households who reported log changes in consumption and house value larger than 1; this corresponds to increases of about 170 percent and decreases of about 63 percent. Such huge values in the left- and right-hand side variables can distort OLS estimates of a linear relationship.⁸

⁸It is worth noting that the observations with the largest house value growth (decline) also have the largest consumption growth (decline); these outliers are especially dangerous because of this positive relationship.

TABLE A.1: Choosing Stable Families

Selection Criterion	Observations		Households
	Dropped	Remaining	
Initial sample		113,659	10,017
Neither owned nor rented	5,402	108,257	9,858
Unmarried	36,285	71,972	6,125
Head aged 24 or younger	4,666	67,306	5,726
Zero/invalid family income	28	67,278	5,725
Lived abroad	188	67,090	5,721
Zero food consumption	252	66,838	5,713
Invalid housing information ^a	3,748	63,090	5,335
Fewer than 5 obs. in 1968–1993 ^b	6,195	56,895	3,312
Not in sample 6 consecutive years	11,776	45,119	2,960
<i>Eliminate early release observations</i>			
Year \geq 1994	4,983	40,136	2,960
Fewer than 6 observations	494	39,642	2,832

NOTE. Table describes the selection of the base sample of stable households. The main criteria for inclusion are that the household consist of a married couple, have a head older than 24, and respond to at least six consecutive PSID questionnaires.

^aFor homeowners, a reported house value of 0 or 999,999 (an invalid code), or less than \$1,000. For renters, annual reported rent that is missing, zero, or greater than \$60,000; in 1988 and 1989 rent information was not collected, these observations were not deleted.

^bNote that this requirement knocks out many of the extra sample of Latino families introduced in the early 1990s.

TABLE A.2: Sample Selection Within Stable Households

Criterion	Observations	
	Dropped	Remaining
<i>Base restrictions</i>		
Bad Interview Year ^a	10,249	29,393
Missing income information	1,813	27,580
Renters (previous year)	2,140	25,440
Renters (interview year)	116	25,324
<i>Additional restrictions</i>		
Moved in past year	1,246	24,078
New husband/wife	338	23,740
Change in geocode for State ^b	48	23,692
Family composition change ^c	4	23,688
Family size changed	3,344	20,344
Adults \neq 2	236	20,108
Bottom 1% of food consumption growth	202	19,906
Top 1% of food consumption growth	202	19,704
Bottom 1% of house value growth	193	19,511
Top 1% of house value growth	195	19,316

NOTE. The table shows the base selection and the additionally restricted selection of household-years from the sample of stable families. The base selection criteria are the minimum necessary criteria to run the regression of interest; the additional restrictions eliminate outliers and households that might change consumption for reasons unrelated to house value appreciation.

^aYears in which consumption growth cannot be determined: 1968, 1969, 1973, 1974, 1988, 1989, and 1990

^bIn principle, households that did not move should *ipso facto* not change States; most of these observations were changes in invalid/missing State geocodes. Rather than impute other values, I choose to simply drop these observations.

^cOther than new husband/wife or new children

B Sample Statistics

Figures B.1 and B.2 present sample statistics of the income and consumption of households, respectively. Note that households used in the sample generally have higher incomes and consume more than the broader sample of stable families, especially in the early years of the dataset. Recall that to construct growth rates of consumption and house values for each family we must discard the first observation. Thus families in the sample are generally older, more experienced and so on. Also notice that reported food expenditures at home have been declining while food expenditures away from home have been increasing; this is probably related to the increasing female labor force participation rates over the sample. Figure B.3 shows the mean and median reported house value for both groups; again, the levels are slightly different (especially in the early years) but the paths are similar.

Figure B.4 presents the distribution of the observed growth in house value and total food consumption; the figures can be thought of as smoothed histograms. As can be seen on the figures, the one-year reported changes in both food expenditures and house values tend to be quite large, and vary a great deal.

TABLE B.1: Sample Means

Variable	All	Age Ranges				
		25–34	35–42	43–51	52–62	63–95
Family Size	3.38 (0.80)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)
Age, Head	48.03 (14.44)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)
Age, Wife	45.04 (14.14)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)
Hours. Head ^a	1.85 (1.03)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)
Hours. Wife ^b	0.82 (0.89)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)	0.00 (0.00)

^aThousands of hours worked per year.

^bThousands of hours worked per year.

FIGURE B.1: Sample Income Statistics

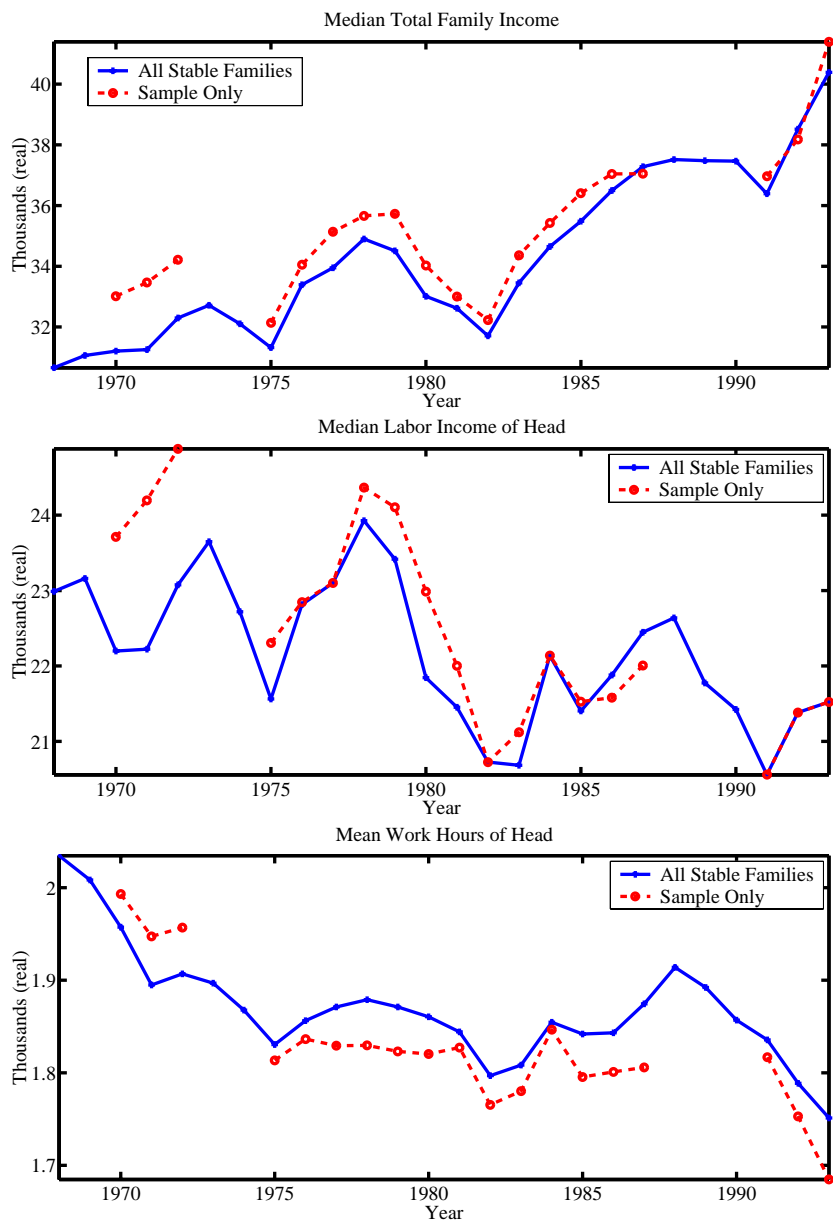


FIGURE B.2: Sample Consumption Statistics




FIGURE B.3: House Values


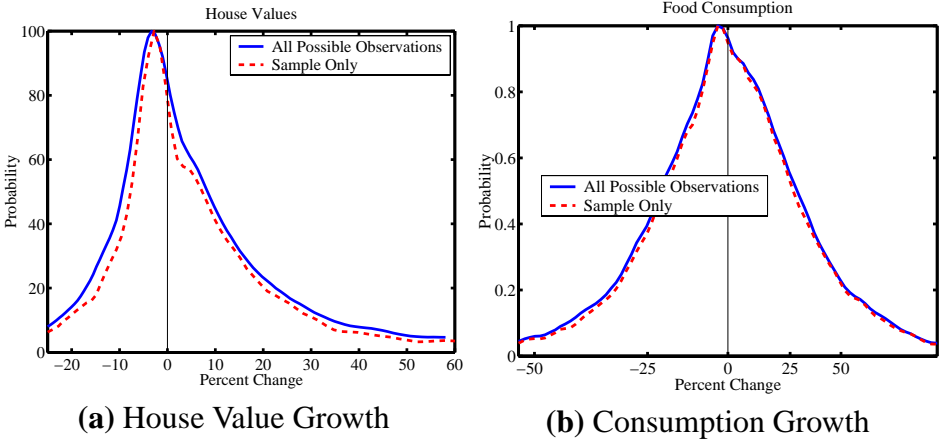


FIGURE B.4: Densities



C Alternate Specifications

I now present estimated elasticities under every possible combination of dataset and specification. In addition, I report results for different age groups; I divided the sample into two groups at the median and into terciles.

The main deviation from the results presented in the main body of the paper is that, when using the unrestricted or full sample described in table A.1 (that is, the larger sample before the additional restrictions of table A.2 have been imposed), the consumption elasticity for the lowest age quintile drops to zero. The reason for this is that households in the lowest age quintile are the likeliest to experience changes in family composition and to move. As I argued in appendix A households that move or that change composition are likely to change their reported food consumption for reasons unrelated to housing wealth, and in ways that are quite difficult to control.

The other main insight is that, by breaking the sample into terciles, it is the middle tercile that has the largest elasticity, the oldest tercile that has the second-largest elasticity and the youngest tercile that has the lowest elasticity. In all cases, though, the estimated elasticities are statistically different from zero.

Tables C.1 and C.2 give results using the unrestricted (i.e. the “full”) sample of stable households, described in table A.1. The tables present the estimated elasticities under all six different specifications for all ages, ages divided by the median, into terciles and into quintiles.

Tables C.3 and C.4 give results using the restricted sample of stable households, described in table A.2. The tables present the estimated elasticities under all six different specifications for all ages, ages divided by the median, into terciles and into quintiles.

TABLE C.1: Consumption Elasticities Using Full Sample

All	Age Quintiles				
	25–34	35–42	43–51	52–62	63–95
<i>Base</i>					
0.0276 (0.0100)	0.0097 (0.0216)	0.0072 (0.0232)	0.0015 (0.0241)	0.0757 (0.0223)	0.0222 (0.0225)
<i>State Dummies</i>					
0.0274 (0.0100)	0.0071 (0.0217)	0.0103 (0.0234)	0.0024 (0.0243)	0.0746 (0.0224)	0.0217 (0.0226)
<i>State Dummies and State House Price Gains</i>					
0.0307 (0.0115)	0.0210 (0.0245)	0.0042 (0.0259)	−0.0054 (0.0291)	0.0804 (0.0264)	0.0306 (0.0260)
<i>Permanent Income Controls</i>					
0.0276 (0.0100)	0.0104 (0.0216)	0.0072 (0.0232)	0.0016 (0.0241)	0.0759 (0.0223)	0.0224 (0.0225)
<i>Permanent Income Controls and State Dummies</i>					
0.0274 (0.0100)	0.0077 (0.0217)	0.0103 (0.0234)	0.0027 (0.0243)	0.0746 (0.0224)	0.0218 (0.0226)
<i>Permanent Income Controls + State Dummies + State House Price Gains</i>					
0.0307 (0.0115)	0.0219 (0.0245)	0.0042 (0.0260)	−0.0055 (0.0291)	0.0804 (0.0264)	0.0307 (0.0260)

NOTE. Table gives OLS estimates of the consumption elasticity of housing wealth gains using the full sample; standard errors are in parentheses. Results are reported for all ages and for each age quintile separately.

TABLE C.2: Consumption Elasticities Using the Full Sample

	Age Ranges				
	Median		Terciles		
	25–46	47–95	25–39	40–54	55–95
<i>Base</i>	–0.0001 (0.0145)	0.0460 (0.0140)	0.0125 (0.0167)	0.0393 (0.0180)	0.0276 (0.0177)
<i>State Dummies</i>	–0.0002 (0.0145)	0.0458 (0.0140)	0.0111 (0.0168)	0.0404 (0.0181)	0.0275 (0.0178)
<i>State Dummies and State House Price Gains</i>	0.0045 (0.0164)	0.0495 (0.0164)	0.0199 (0.0187)	0.0401 (0.0214)	0.0279 (0.0205)
<i>Permanent Income Controls</i>	–0.0001 (0.0145)	0.0459 (0.0140)	0.0128 (0.0167)	0.0394 (0.0180)	0.0276 (0.0177)
<i>Permanent Income Controls and State Dummies</i>	–0.0002 (0.0145)	0.0458 (0.0140)	0.0114 (0.0168)	0.0407 (0.0181)	0.0275 (0.0178)
<i>Permanent Income Controls + State Dummies + State House Price Gains</i>	0.0045 (0.0164)	0.0495 (0.0164)	0.0203 (0.0187)	0.0399 (0.0214)	0.0279 (0.0205)

NOTE. Table gives OLS estimates of the consumption elasticity of housing wealth gains using the unrestricted (or full) sample; standard errors are in parentheses. Results are reported for ages above and below the median and for each age tercile separately.

TABLE C.3: Consumption Elasticities Using Restricted Sample

All	Age Quintiles				
	25–34	35–42	43–51	52–62	63–95
<i>Base</i>					
0.0394 (0.0105)	0.0419 (0.0251)	0.0084 (0.0247)	0.0283 (0.0263)	0.0734 (0.0220)	0.0305 (0.0216)
<i>State Dummies</i>					
0.0390 (0.0105)	0.0373 (0.0254)	0.0089 (0.0249)	0.0295 (0.0265)	0.0718 (0.0223)	0.0281 (0.0217)
<i>State Dummies and State House Price Gains</i>					
0.0471 (0.0118)	0.0448 (0.0280)	0.0136 (0.0280)	0.0361 (0.0310)	0.0915 (0.0255)	0.0277 (0.0238)
<i>Permanent Income Controls</i>					
0.0394 (0.0105)	0.0419 (0.0251)	0.0091 (0.0247)	0.0284 (0.0263)	0.0732 (0.0220)	0.0308 (0.0216)
<i>Permanent Income Controls and State Dummies</i>					
0.0390 (0.0105)	0.0373 (0.0254)	0.0094 (0.0249)	0.0296 (0.0265)	0.0713 (0.0223)	0.0282 (0.0217)
<i>Permanent Income Controls + State Dummies + State House Price Gains</i>					
0.0472 (0.0118)	0.0450 (0.0280)	0.0145 (0.0280)	0.0359 (0.0310)	0.0916 (0.0255)	0.0282 (0.0238)

NOTE. Table gives OLS estimates of the consumption elasticity of housing wealth gains using the restricted sample; standard errors are in parentheses. Results are reported for all ages and for each age quintile separately.

TABLE C.4: Consumption Elasticities Using Restricted Sample

	Age Ranges				
	Median		Terciles		
	25–46	47–95	25–39	40–54	55–95
<i>Base</i>	0.0237 (0.0160)	0.0490 (0.0140)	0.0135 (0.0191)	0.0677 (0.0190)	0.0366 (0.0171)
<i>State Dummies</i>	0.0238 (0.0161)	0.0483 (0.0141)	0.0104 (0.0192)	0.0702 (0.0190)	0.0360 (0.0172)
<i>State Dummies and State House Price Gains</i>	0.0320 (0.0181)	0.0561 (0.0158)	0.0152 (0.0211)	0.0925 (0.0223)	0.0358 (0.0190)
<i>Permanent Income Controls</i>	0.0238 (0.0160)	0.0490 (0.0140)	0.0138 (0.0191)	0.0678 (0.0190)	0.0366 (0.0171)
<i>Permanent Income Controls and State Dummies</i>	0.0238 (0.0161)	0.0482 (0.0141)	0.0107 (0.0192)	0.0704 (0.0191)	0.0359 (0.0172)
<i>Permanent Income Controls + State Dummies + State House Price Gains</i>	0.0323 (0.0181)	0.0564 (0.0158)	0.0159 (0.0212)	0.0926 (0.0223)	0.0361 (0.0191)

NOTE. Table gives OLS estimates of the consumption elasticity of housing wealth gains using the restricted sample; standard errors are in parentheses. Results are reported for ages above and below the median and for each age tercile separately.