

House Prices and Consumer Spending *

David Berger
Northwestern University

Veronica Guerrieri
University of Chicago

Guido Lorenzoni
Northwestern University

Joseph Vavra
University of Chicago

June 9, 2015

*Berger: Northwestern University, Department of Economics, 2001 Sheridan Road, Evanston, IL, 60208 (e-mail: david.berger@northwestern.edu); Guerrieri: University of Chicago, Booth School of Business, 5807 South Woodlawn Avenue, Chicago, IL, 60637 (e-mail: vguerrie@chicagobooth.edu); Lorenzoni: Northwestern University, Department of Economics, 2001 Sheridan Road, Evanston, IL, 60208 (e-mail: guido.lorenzoni@northwestern.edu); Vavra: University of Chicago, Booth School of Business, 5807 South Woodlawn Avenue, Chicago, IL, 60637 (e-mail: joseph.vavra@chicagobooth.edu). We thank Sasha Indarte for excellent research assistance.

Abstract

We use a life-cycle model with income uncertainty and incomplete markets to derive a new theoretical benchmark for the response of consumption to house price movements. In a variety of specifications of our model the individual response of consumption to a one-time, permanent house price change is well approximated by a simple formula which can be measured in the data: the marginal propensity to consume out of temporary income times housing holdings. In a baseline calibration, with perfectly liquid housing and no rentals, the model generates very large housing wealth effects, as households with low net worth buy relatively large houses using leverage and so combine high MPCs and high housing holdings. We then extend the model to allow for illiquid housing, the option to rent, and more realistic mortgages. Housing wealth effects are still large, but smaller than in the baseline, as households with low net worth self-select into renting. The effects in the extended model are in the range of recent empirical estimates. We also use the model to explore the consumption implications of a boom-bust in house prices. To generate a boom-bust in residential investment, we introduce time-varying expectations of future housing appreciation. The endogenous evolution of household debt plays a crucial role in consumption dynamics: increased leverage in the boom contributes substantially to the consumption contraction in the bust.

1 Introduction

The last recession in the US has been characterized by a surprising large and persistent weakness in consumer spending. Many observers have pointed to the evolution of house prices and household debt as important factors in explaining consumption dynamics before and during the recession. More generally, an important open question is what are the effects of boom-bust cycles in house prices and household debt on consumer spending. A large and growing empirical literature argues that consumption responds strongly to house price movements, thus suggesting a large role for housing wealth in consumption dynamics. The theoretical rationale for these “housing wealth effects” is less clear. The standard permanent-income model implies that wealth effects should be very small. A commonly held view claims that housing wealth effects should be especially small because most households own housing assets but are not net traders of them. What kind of models can deliver the large effects found in the empirical literature? What are the mechanisms at work? Should we expect housing wealth effects to be larger when the household sector is more levered? Is it important that housing can be used as collateral? Does it matter what underlying shocks are driving the boom-bust in house prices?

In this paper, we study these questions in the context of an incomplete market model with income uncertainty and housing that serves as collateral. Our main messages are as follows: (i) A fairly standard model, reasonably calibrated, can produce quite large responses of consumption to house prices; (ii) These responses can be as large as in the empirical literature; (iii) The economy’s initial conditions, and in particular the joint distribution of housing and debt, are very important in determining the size of these responses.

The paper offers theoretical results that help to inspect the mechanism behind these findings. In our baseline model, the individual response of consumption to a one-time, permanent house price change is given by a simple formula: the individual marginal propensity to consume out of temporary income shocks (MPC) times individual housing holdings. The aggregate response is then determined by the joint distribution of the MPC and housing. Both the MPC and housing are endogenous in the model and their joint distribution in the population varies in different specifications of the model and following different shocks. So the usefulness of the formula is that it helps to explain why different specifications generate different average responses. It also helps to explain why, following different histories of past

shocks, which produce different initial conditions, an economy displays different aggregate responses.

The formula also holds for a special case with no income uncertainty which satisfies the permanent-income hypothesis and allows us to decompose the total house price effect into substitution and income effects, a collateral effect, and an endowment effect. These exercises help to clarify the connection between our results and the small-wealth-effects views described above.

Our baseline is a partial equilibrium model calibrated to life-cycle wealth data from the Survey of Consumer Finances. The aggregate effect of an unexpected, permanent house price shock gives an elasticity in the range of 0.34-0.40, depending on the specification. To explain this result notice first that in a model with income uncertainty and precautionary savings, the consumption function is concave in net worth and so the MPC can be large for low net worth agents. Moreover, in a model where housing services enter the utility function, holdings of housing wealth are not proportional to net worth. In particular, low net-worth agents choose optimally to borrow and hold housing that is a multiple of net worth. Therefore, the model generates a joint distribution with a sizeable mass of agents at high housing levels and, at the same time, low net worth and high MPC. Using our formula then implies that the aggregate response can be large.

The baseline model features no transaction costs in housing and no rental option. Therefore, it produces unrealistic volumes of house trading and cannot match the large fractions of renters in the population. We then extend the model to allow for illiquid housing and for a rental option and re-calibrate the model to match realistic levels of house trading and a realistic fraction of renters. In the extended model, our simple formula does not hold exactly, but still provides a good approximation for realistic choices of parameters. Somewhat surprisingly, the presence of adjustment costs in housing has only minor quantitative implications for the effect of house prices on consumption. The main quantitative difference between the baseline model and the extended model comes from the presence of renters. The presence of renters not only means that some agents do not own housing, and thus have a zero endowment effect, it also means that agents choose optimally to select into renting and the agents who choose to do so are precisely the low net-worth, high MPC agents that generate the biggest responses. Therefore, the aggregate response is smaller in the model with renters, around 0.21-0.24. These numbers are in

line with recent empirical estimates of consumption responses to house prices.

We also take our simple formula directly to the data to use it as a sufficient statistic for the aggregate response of consumption. The main challenge here is to estimate MPC conditional on housing holdings. We do so extending the approach of Blundell, Pistaferri and Preston (2008) and obtain numbers that are similar to those obtained by simulation.

Finally, we use the model to explore the consumption implications of a boom-bust in house prices. To do so, requires taking an extra step relative to the simple partial equilibrium exercises conducted in the rest of the paper. In particular, the behavior of residential investment in boom-bust cycles suggests that the driving force behind the increase in house prices must be a housing demand shock, which increases the demand for housing at the same time as prices are increasing. To produce a demand shock, we introduce time-varying expectations of future house price appreciation. Doing so means that the consumption response will be driven by two forces: the house price effect—which is the main subject of this paper—and the effect of changes in expected appreciation. Once we account for both effects, we still obtain fairly large responses of consumption. But the most important lesson from this exercise is that the endogenous evolution of household debt plays a crucial role in explaining consumption dynamics. Increased leverage in the boom contributes substantially to generating a sizeable consumption contraction in the bust.

2 Literature Review

There is a large empirical and theoretical literature studying the effect of housing prices on nondurable consumption. While there are many ways of measuring the consumption wealth effect, whenever possible we will discuss estimates in elasticity terms - that is how much in percentage terms does nondurable consumption change when log housing prices change by one percent - because our theory speaks directly to the magnitude of this elasticity.

Roughly speaking, there are two types of empirical studies: 1) those using aggregate time-series evidence and 2) those using household level micro data. Two of the most widely cited studies using aggregate data are Case, Quigley and Shiller (2013) and Carroll, Otsuka and Slacalek (2011). Case, Quigley and Shiller (2013) using cross-state variation and find that the elasticity of nondurable consumption

ranges from 0.03 and 0.18 with most of the their estimates centered around 0.10. These numbers are also consistent with an earlier study done by these authors (Case, Quigley and Shiller, 2005) which used a panel of developed countries. Carroll, Otsuka and Slacalek (2011) use a time-series based method that exploits the sluggishness of consumption growth to distinguish between immediate and eventual wealth effects. Using U.S. data, they estimate that the immediate (next-quarter) marginal propensity to consume from a \$1 change in housing wealth is about 2 cents, with an eventual effect around 9 cents.

There have also been a number of recent studies that exploit household level micro data. A key advantage of micro data is that there is much more variation in consumption and housing prices, however, there is also significantly more measurement error and often a lack of desired covariates that would aid in the interpretation of these micro regressions. Most importantly, to isolate a “pure” housing wealth effect, one would need data on expenditure of individual households before and after some truly exogenous change in his home values and these truly exogenous events are difficult to find. Perhaps reflecting these difficulties, the first generation of these studies produced estimates that varied widely, despite using similar datasets. Campbell and Cocco (2006) use British data from the U.K. Family Expenditure Survey and find an average elasticity of nondurable consumption to house prices of about 1.2 and find larger results for old homeowners than for young homeowners. Attanasio, Blow, Hamilton, and Leicester (2008), in contrast, find that the average nondurable consumption elasticity is 0.15, but that that consumption response of young homeowners is much stronger (0.21) than it is for old homeowners (0.04) despite using a very similar data set.

The study that perhaps comes closest to the ideal micro experiment is Mian, Sufi and Rao (2013). In this study, they exploit the Saiz (2010) housing supply instrument to isolate plausibly exogenous changes in housing prices and then examine the effect of this exogenous housing price variation on zip and county level measures of expenditures. Their baseline estimates of the nondurable consumption elasticity are between 0.13 and 0.26.¹ We use the Mian, Sufi and Rao (2013) estimates as our benchmark for two reasons. First, they are best identified estimates in the micro

¹They report estimates between 0.5-0.8. However, given their methodology these estimates need to be scaled by housing wealth/total wealth to be comparable to the other estimates listed above. Since the mean housing wealth to total wealth ratio in their data is between 0.25-0.33, this implies elasticity estimates ranging from 0.13 to 0.26.

literature. Second, they exploit cross-sectional evidence and thus abstract from aggregate GE effects, their estimates are most directly comparable to our model results since in the vast majority of theoretical exercises we take housing prices as given.

The extrant theoretical literature also exhibits substantial heterogeneity in terms of proposed channels. Much of this reflects that there is significant heterogeneity in how housing price changes will affect various agents. Indeed, some households will be better off when home prices rise and others will be worse off. Renters are unambiguously worse off when the cost of housing consumption increases because now they face a higher cost of shelter without bearing any of the upside. The impact of home price appreciation on homeowners is less clear. Standard economic theory offers a framework for thinking about wealth effects. According to the permanent income hypothesis, households consume a constant fraction of the present discounted value of their lifetime resources. Changes in wealth that permanently alter households' resources therefore should cause consumption to change in the same direction. However, even in the PIH world there are many channels through which housing prices could affect nondurable consumption and even more so when one allows for collateralized borrowing. As will become clear in the next few sections there are (at least) four channels through which housing prices change effect consumption: the substitution effect, the income effect, the endowment effect and a collateral effect.

Much of theoretical literature focuses on a subset of these four channels and a main contribution of this paper is to provide a useful taxonomy of these theoretical channels as well as how they related to each other. For example, there is a large literature examining whether housing is really wealth at all. As Sinai and Souleles (2004) point out, everyone is born short housing. Then since households consume housing services and if a homeowner who expects to live in his current house for a long time, a higher house price is simply compensation for a higher implicit rental cost of living in the house. This suggests that absent substitution and collateral effects, the housing wealth effect should be close to zero as the income and endowment effect cancel. We will show that while this effect is present in our model, an alternative decomposition exists which shows that the income, substitution and collateral effects all cancel and all that is left is the endowment effect. This decomposition is the heart of our paper.

Finally, we will show that despite the fact that many pairs or triplets of these

effects cancel, this does not imply that each of the individual effects or the aggregate elasticity is small. In particular, consistent with the evidence in Cooper (2013) and DuFusco (2014) that the collateral channel is large. Similarly, consistent with the results in Mian, Sufi and Rao (2013) we find that the endowment effect is also substantial.

3 Model

We consider an incomplete market model with heterogeneous agents facing idiosyncratic risk a la Bewley, enriched in two main dimensions. First, we introduce housing in the model, which both gives utility to the households and represents an alternative asset in their saving decisions. Second, we introduce a life-cycle component that we believe is important to think about housing and saving decisions.

Time is discrete and run forever. The economy is populated by a continuum of households indexed by i , who live J periods. The first J_y period correspond to working age, the next J_o periods correspond to retirement.

Households have access to two types of assets: a risk-free asset and housing. We denote the holdings of the two assets, respectively, A_{it} and H_{it} . The risk-free asset is perfectly liquid and yields a constant interest rate r . The housing stock yields housing services one for one to the household, depreciates at rate δ , and has price P_t which is taken as given. If household i decides to trade housing at time t , it incurs an adjustment cost that is proportional to the value of the house sold

$$\kappa_{it} = FP_t H_{it-1} \mathbb{1}_{H_{it} \neq H_{it-1}},$$

where $\mathbb{1}$ is an indicator function equal to 1 iff $H_{it} \neq H_{it-1}$.

Households maximize expected utility, have time-separable preferences and discount the future at the rate β . The per-period utility is

$$U(C_{it}, H_{it}) = \frac{(C_{it}^\alpha H_{it}^{1-\alpha})^{1-\sigma}}{1-\sigma},$$

where C_{it} is consumption of non-durable goods and the housing stock H_{it} delivers housing services 1:1. The housing and non-housing wealth left by the household at the end of period J go to their offsprings in the next period and enters the household

utility function through the bequest function

$$\frac{\Psi}{1 - \sigma} \left(\frac{\Gamma_i + (1 - \delta - F)P_t H_{it-1} + (1 + r)A_{it-1}}{\hat{P}_{Xt}} \right)^{1 - \sigma},$$

where Γ_i captures the human capital of the offsprings and \hat{P}_{Xt} is a price index that will be defined in the next section, which converts nominal wealth into real wealth.

Households face an exogenous income process. When the household is of working age, the income is

$$Y_{it} = \exp\{\chi(j) + z_{it}\},$$

where $\chi(j)$ is an age-dependent parameter, j is the age of the household at time t , and z_{it} is a transitory shock that follows an AR1 process

$$z_{it} = \rho z_{it-1} + \varepsilon_{it}.$$

When the household is retired, income is given by a social security transfer, which is a function of income in the last working-age period. The social security transfer is modeled as in Guvenen and Smith (2014).

The household's budget constraint is

$$C_{it} + P_t(H_{it} - (1 - \delta)H_{it-1}) + A_{it} = Y_{it} + (1 + r)A_{it-1} - \kappa_{it}.$$

Moreover, households face the borrowing constraint

$$-(1 + r)A_{it} \leq (1 - \theta)(1 - \delta)P_{t+1}H_{it}, \quad (1)$$

that is, their debt position at the beginning of $t + 1$ cannot be larger than a fraction $1 - \theta$ of the value of their house at that time.

4 The Permanent-Income Case

As a preliminary step, it is useful to consider a special case with a deterministic income path ($\sigma_\varepsilon = 0$), no adjustment costs ($F = 0$), no borrowing constraint, and a constant house price $P_t = P$. Under some additional assumptions, this special case can be solved analytically and gives us a reference model with housing wealth

where the permanent-income hypothesis applies. For simplicity, we also set the human capital of offsprings Γ_i to zero.

The following assumptions on β and Ψ imply perfect consumption smoothing, that is, households consume a constant amount of non-durable goods and housing throughout their lifetime. Namely, assume

$$\beta = (1 + r)^{-1}, \quad \Psi = (1 - \beta)^{-\sigma} \hat{P}_X^{\sigma-1}.$$

Define total spending per period as

$$X_{it} = C_{it} + \left(P - \frac{1 - \delta}{1 + r} P \right) H_{it}.$$

The second term on the right-hand side is current spending on housing services, since the term in parenthesis corresponds to the user cost of housing (i.e., the implicit rental cost). Given our assumptions on β and Ψ the household satisfies the permanent-income hypothesis by keeping spending X_{it} constant over time and equal to the interest on its total wealth, which includes human wealth, housing wealth, and financial wealth:

$$X_{it} = \frac{r}{1 + r} \left[\sum_{\tau=t}^{t+J-j} (1 + r)^{-\tau} Y_{it+\tau} + (1 - \delta) P_t H_{it-1} + (1 + r) A_{it-1} \right].$$

Given X_{it} , Cobb-Douglas preferences imply that a fraction α is allocated to non-durable consumption. Therefore, non-durable consumption is proportional to total wealth. Now consider the effect of an unexpected, permanent shock to the house price. The elasticity of consumption to this shock is equal to the housing share out of total wealth, that is,

$$\frac{dC_{it}/C_{it}}{dP/P} = \frac{(1 - \delta) P H_{it-1}}{\sum_{\tau=t}^{t+J-j} (1 + r)^{-\tau} Y_{it+\tau} + (1 - \delta) P_t H_{it-1} + (1 + r) A_{it-1}}. \quad (2)$$

Since total wealth includes human wealth, under reasonable parametrizations, a permanent-income model typically implies small consumption elasticities to house prices. Notice that a permanent-income model aggregates exactly, so we can look at the elasticity of aggregate consumption to an aggregate shock to house prices. Let us use the 2001 Survey of Consumer Finances (SCF) for some back-of-the-envelope cal-

culations, since we will use the same dataset for our main calibration of households' balance sheets. Namely, matching the corresponding quantities in the 2001 SCF, set $(1 - \delta)PH = 2.15Y$ and $A = -.32Y$, where H is average housing, A average liquid wealth net of debt, and Y average earnings. Using an interest rate of $r = 2.5\%$ and using an infinite horizon approximation, human wealth is computed as $Y/r = \dots Y$. Then the aggregate elasticity of non-durable consumption implied by the model is 0.0514, a small number. Moreover, this elasticity is not sensitive to the level of indebtedness of the household sector. Suppose that, all else equal, household debt goes up by $0.5Y$, a large increase, and re-compute the elasticity at $A = -0.82Y$. We obtain an elasticity of 0.0520, virtually identical to the one at $A = -.32Y$. Finally, notice that if we look at elasticities by age group, elasticities will tend to be higher for old households, since they have a smaller fraction of human wealth over total wealth.

The take out of this section is that a standard permanent-income model produces small elasticities, that these elasticities are minimally affected by the level of the households' debt, and that elasticities tend to be smaller for younger households.

What is the mechanism behind the consumption response in the permanent-income model? We can do a standard decomposition into three effects: a substitution effect, an income effect, and an endowment effect.² It is then possible to look at equation (2) from two angles. First, we can notice that due to the Cobb-Douglas assumption, the income and substitution effect exactly cancel in our model. So the effect in (2) can be interpreted purely as an endowment effect. For an alternative interpretation, notice that in the PIH the consumption of housing services is constant over time, so if we take a household at any point in time, after its first period of life, the effect of the price increase on the value of its endowment and the effect on the net present value of its expenditures on housing services will be roughly equal.³ Therefore, the effect in (2) can also be interpreted as an (almost) pure substitution effect, with the income and endowment effect canceling each other. The last interpretation is consistent with the widely held view that housing wealth effects must

²A permanent change in P is equivalent to a change in the service cost of housing in all future periods. The substitution effect is the shift from housing services in all future periods towards current consumption for a constant present value of expenditure. The income effect is the change in current consumption due to a reduction in the present value of expenditure equal to the increase cost of housing in all future periods. The endowment effect is the change in current consumption due to an increase in the present value of expenditure equal to the increase in the value of the initial housing stock.

³The effects are not exactly equal due to depreciation δ . When $\delta = 0$ they are exactly equal.

be small, because an increase in the price of housing increases the value of an asset for which many people are roughly at zero net trade. It is useful to remark that at this stage, both interpretations of (2) are correct. The first interpretation will be especially useful in what follows, because, as we shall see, it will survive in richer versions of the model.

5 Baseline Model: Precautionary Savings and Borrowing Constraints

Let us now go back to the general model and re-introduce income uncertainty and the borrowing constraint (1). In this section we maintain the assumption that housing is perfectly liquid, that is, there are no adjustment costs for housing, so $F = 0$. In this setup, we derive our main analytical result: the individual consumption response to a permanent change in the house price is given by a simple formula, the product of the individual marginal propensity to consume out of temporary income shocks and the beginning-of-period housing stock.

To set the stage for the result, let us represent the household problem in a recursive way. In order to do so, it is useful to recognize that the only individual state variables for household i at time t are total wealth $W_{it} \equiv (1 - \delta)P_{t+1}H_{it-1} + (1 + r)A_{it-1}$, the idiosyncratic income variable z_{it} , and age j_{it} . In recursive notation, we can then define the household value function in terms of W and $s \equiv (z, j)$ as follows

$$V_t(W, s) = \max_{C, H, A, W'} U(C, H) + \beta E [V_{t+1}(W', s')]$$

subject to

$$C + P_t H + A = Y(s) + W$$

$$W' = (1 - \delta) P_{t+1} H + (1 + r) A$$

$$(1 - \theta) (1 - \delta) P_{t+1} H + (1 + r) A \geq 0,$$

where $Y(s) = \exp\{\chi(j) + z\}$, $z' = \rho z + \varepsilon'$ and $j' = j + 1$. To complete the description of the problem the bequest motive gives us the terminal condition:

$$V_t(W, s) = \frac{\Psi}{1 - \sigma} \left(\frac{\Gamma(s) + W}{\hat{P}_{Xt}} \right)^{1 - \sigma},$$

for all $s = (j, z)$ with $j = J + 1$.

In a model with constant house prices, the user cost of housing services (i.e., the implicit rental rate) is also constant and equal to $P(r + \delta)/(1 + r)$. Since the non-durable good is the numeraire, the CPI is then

$$P_X = \alpha^{-\alpha} (1 - \alpha)^{-(1-\alpha)} \left(P \frac{r + \delta}{1 + r} \right)^{1-\alpha},$$

and we can set $\hat{P}_X = P_X$ in the bequest function.

We are now ready to prove our main analytical result.

Proposition 1 *The individual response of nondurable consumption to a permanent, unexpected change in house prices, in the model with no adjustment costs, can be written as*

$$MPC_{it} \cdot (1 - \delta) H_{it-1}, \quad (3)$$

where MPC_{it} is the individual marginal propensity to consume out of transitory income shocks. Hence, the individual elasticity of nondurable consumption to house prices is

$$\eta_{it} = MPC_{it} \cdot \frac{(1 - \delta) P_t H_{it-1}}{C_{it}}.$$

Proof. With constant prices the Bellman equation can be written in the form

$$V(W, s) = \max_{C, H, W'} U(C, H) + \beta E [V(W', s')],$$

$$W' = (1 + r) \left[W + Y(s) - C - \frac{r + \delta}{1 + r} PH \right],$$

$$W' \geq \theta (1 - \delta) PH.$$

Given Cobb-Douglas/CRRA preferences this can be rewritten in terms of the value of housing $\tilde{H} = PH$ as

$$V(W, s) = \max P^{-(1-\sigma)(1-\alpha)} U(C, \tilde{H}) + \beta E [V(W', s')],$$

$$W' = (1 + r) \left[W + Y(s) - C - \frac{r + \delta}{1 + r} \tilde{H} \right],$$

$$W' \geq \theta (1 - \delta) \tilde{H}.$$

Notice that the expression $P^{-(1-\sigma)(1-\alpha)}$ is a constant that multiplies the utility function in each period. It also appears in the same form in the bequest function, through the CPI term \hat{P}_X , as derived in the text. We conclude that the policy function $C(W, s)$ for the problem above is independent of the level of P . Therefore, the response of consumption of household i to a permanent shock to P at t is

$$\frac{\partial C(W_{it}, s_{it})}{\partial W} (1 - \delta) H_{it-1}.$$

To complete the argument notice that $\partial C(W_{it}, s_{it}) / \partial W$ is equal to MPC_{it} . ■

This proposition shows that the individual consumption response to a permanent change in house prices can be represented by a simple formula. Notice that both objects in the formula are endogenous, so the Proposition does not give a closed form expression for the response. However, the formula will be very useful to understand how endogenous forces determine whether an economy is more or less sensitive to a house price shock. In particular, house price shocks will have a bigger effect if there are enough high-MPC consumers in the economy with sizeable holdings of housing wealth. The usefulness of the formula comes from the fact that it will hold approximately also in richer versions of the model with adjustment costs and with the option to rent, as we shall see in Section If one can get reliable estimates of MPCs conditional on housing wealth, the formula can also be used as a sufficient statistic to estimate the model-implied response of consumption to house price shocks, without the need to structurally estimate the full model. We will explore this approach in Section

At first sight, formula (3) may appear obvious or tautological.⁴ The non-obvious content of the formula lies in the fact that the MPC is the marginal propensity to consume out of temporary income shocks. To further grasp the non-obvious nature of the result, it is useful to think of the several effects that determine consumer behavior when house prices change. Namely there is: (1) a substitution effect that makes households substitute away from housing, which is now more expensive, towards non-durable consumption; (2) an income effect which makes households feel overall poorer in real terms because housing services are more expensive; (3) a collateral effect which comes from the fact that households can now borrow more given that the value of their collateral has permanently increased; (4) there is an

⁴Of course, the need of a proof may suggest otherwise.

endowment effect which increases the value of the housing stock owned when the shock hits. The upshot of formula (3) is that the first three effects exactly cancel out, so the total effect is equal to the endowment effect. Notice that this result extends the result obtained in the simple permanent-income model. The crucial difference is that the magnitude of the MPC is very different in the two versions of the model. In numerical examples, we will show that although the sum of effects (1) to (3) is zero, each effect in insulation can still be large and we'll discuss interpretations of these effects.

Proposition 1 can be extended to the case of stochastic house prices that follow a random walk with drift and to the case in which additional risky assets are present with i.i.d. stochastic rates of return. If we abandon the assumption of Cobb-Douglas preferences in non-durables and housing services...

5.1 Calibration

We now explain our calibration choices and report the baseline parameters in Table 1.

The model is annual. We interpret the first period of life of a household as age 25. Households work for $J_y = 35$ years (between 25 and 59) and are retired for $J_o = 25$ years (between 60 and 84). We set the interest rate $r = 3\%$. In our baseline calibration we choose an intertemporal elasticity of substitution σ equal to 2, but we also present results for other values of σ . We set the housing depreciation rate δ to 2.2% to match the ratio of depreciation to the housing stock between 1960-2014 in BEA data. We for now work with the frictionless housing adjustment version of the model and so set F to 0.0. The collateral constraint parameter θ represents the mortgage minimum down payment and we set it equal to 0.25.

Let us now move to the income process. Recall the income process during working age has a life-cycle component and a temporary component. The life-cycle component is chosen to fit a quadratic regression of yearly earnings on age from the PSID. The temporary component z follows an AR1 process with autocorrelation $\rho = 0.91$ and a standard deviation $\sigma = 0.21$ to match the autocorrelation and variance of PSID earnings (after taking out the life-cycle component) following Floden and Linde' (2001). For the income process during retirement we assume a social security process as in Guvenen and Smith (2014).

The three remaining parameters to calibrate are: the housing coefficient α , the

Table 1: *Parameter Values*

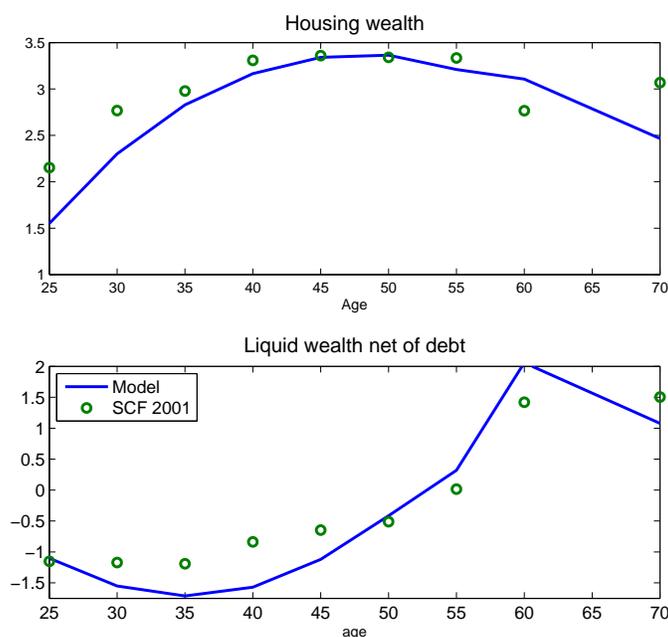
Parameter	Value
σ	2
r	2.4%
δ	2.2%
F	0%
θ	0.25
ρ	0.91
σ	0.21
α	0.853
β	0.9247
Ψ	2,336

Note:

discount factor β , and the bequest parameter Ψ . We choose these parameters jointly to match life-cycle profiles of housing and non-housing wealth in the data. Namely, from the 2001 Survey of Consumer Finances (SCF) we compute average housing wealth and average liquid wealth net of debt for households in 9 age bins (25-29, 30-34, ..., 60-64, 65 and over). We initialize the model by giving agents of age 25 holdings of housing and liquid assets and income to match the distribution of age 23-27 households in the 2001 SCF.⁵ Our model is very stylized for retired agents as we let them face a deterministic income flow and no other source of risk. Therefore, we prefer to focus our calibration and our predictions on working-age agents. For the calibration, we do it by grouping all agents 65 and above in a single bin. Our notion of liquid wealth net of debt includes all assets in SCF excluding retirement accounts for agents before retirement, but including retirement accounts for agents above 60. In the model, we assume that retirement accounts take the form of a lump sum transfer which is made available upon retirement and is a proportion Ξ of the labor income prior to retirement. We choose the parameters α, β, Ψ and Ξ to minimize the quadratic distance between the sequences of housing and liquid wealth by age bin in the data and the corresponding sequences generated by the model. We choose to target liquid wealth in the data rather than networth excluding housing wealth because this delivers MPCs which are more in line with empirical estimates.

⁵Using only age 25 households implies fairly small samples and lots of measurement error. Results are similar when using age 22-25 households but many age 20-22 households are still in school.

Figure 1: Housing and Net Liquid Wealth over the Life Cycle: Data vs Model

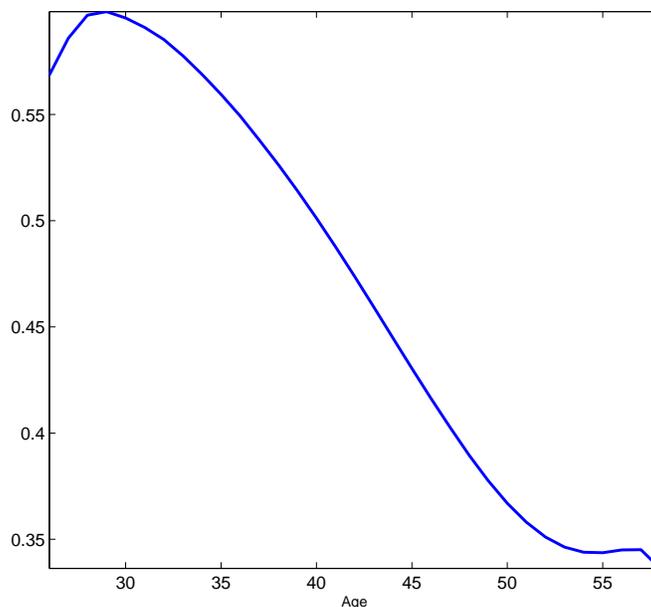


Furthermore, this is consistent with the observation in Kaplan and Violante (2014) that many households are “wealthy-hand-to-mouth.” While it would be desirable to separately model liquid and illiquid wealth in addition to the choice of housing, this would substantially complicate the analysis. The majority of non-housing illiquid wealth is held in retirement accounts which have large penalties for accessing prior to retirement but become fully liquid after retirement. Thus, we believe that our calibration strategy provides a reasonable description of the fraction of wealth which can be easily accessed both prior to and after retirement. Nevertheless, targeting somewhat higher values of wealth does not substantively change any of our conclusions.

Figure 1 shows the fit of the calibrated model in terms of average housing wealth (top panel) and average liquid wealth net of debt (bottom panel), by age bin.⁶ The solid line are the model predictions and the circles are from the 2001 SCF. Despite its simplicity the model delivers an acceptable fit, the main discrepancy being too little housing in the early periods and too much debt in the mid 30s and 40s.

⁶The point for the age bin 65 and over is plotted at age 70.

Figure 2: Elasticities over the Life Cycle: Frictionless Model



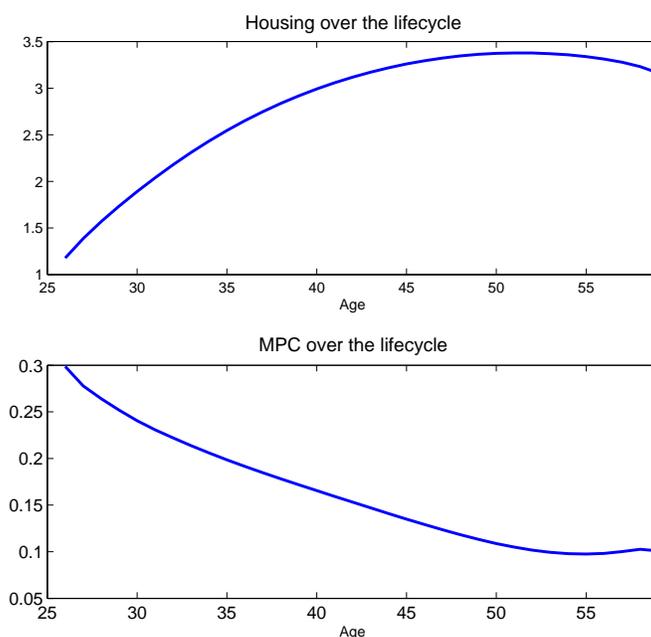
5.2 The effects of a permanent house price shock

We now turn to the model predictions for a simple comparative static exercise on house prices. Namely, we look at a negative, unexpected, permanent reduction in house price by -1% and look at the instantaneous response of non-durable consumption.

The elasticity of average consumption, computed over all age groups, is 0.34. Figure 2 reports the elasticity for different age bins, with values as high as 0.6 for agents in their late 20s. Elasticities are higher for younger agents and decline steadily after 30. The overall take out is that magnitudes are much larger than in the benchmark permanent-income model. We can use our formula (3) to interpret this result.

In Figure 3 we plot separately the two elements of the formula: the MPC and the housing holdings, averaged by age. The reason for the high elasticities of young agents is in the fact that they have a high MPC and, at the same time, have substantial holdings of housing. In a precautionary saving model the MPC depends on total net worth W and is decreasing in W due to the concavity of the consumption function. Young agents own houses but finance them with debt, so they have low net worth in spite of having relatively high housing wealth. This helps to explain

Figure 3: Understanding the Life Cycle: Frictionless Model



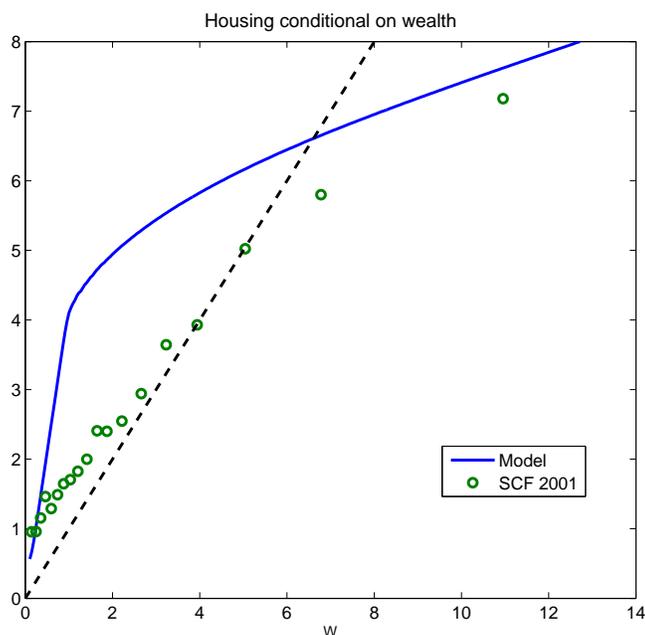
why they have a combination of high MPC and substantial housing wealth.

The discussion above highlights that an important feature of the model, which contributes to deliver high elasticities, is that housing is not a constant fraction of net worth. In Figure 4 we plot the relation between net worth and housing wealth in the data and in the model. Both the model and the data show that housing wealth grows less than proportionally with net work. The relationship starts above the 45 degree line, with net-worth poor, indebted home-owners for which housing wealth is larger than total wealth. Only at relatively high levels of net worth the relation crosses the 45 degree line.

6 Extended Model: Adjustment Costs and Rentals

Proposition 1 shows that in models with frictionless housing adjustment, the consumption response to house price changes depends solely on the product of household MPCs and household housing values. Calibrations of these models which are consistent with household micro data in turn imply very large consumption responses to house price movements as many households borrow substantial amounts

Figure 4: Housing and Net Worth: Data vs Model



to purchase relatively large houses, yet have large MPCs due to this borrowing. The assumption that housing can be adjusted frictionlessly allows us to derive the simple sufficient statistic. Many existing models of consumption and results in the literature use models which satisfy the conditions for Proposition 1 and so can be fully characterized by our result.

Nevertheless, the assumption that housing can be adjusted frictionlessly is clearly counterfactual: brokers' fees and search costs make housing adjustment costly. Berger and Vavra (2015) show that the average annual frequency of housing adjustment in the PSID is just under 5%. Costly refinancing also makes housing equity less than fully liquid. Bhutta and Keys (2014) show that only around 12% of households with positive equity refinance annually to extract that equity. In addition, a large fraction of households rent rather than owning and so are unaffected by the collateral and endowment channel. These renters are also disproportionately young households with lower accumulated wealth and potentially different MPCs. Finally, the results in the previous section assume a symmetric collateral constraint so that when house prices fall and households violate the collateral constraint they must immediately put more cash into their house to meet this constraint. In reality, collateral con-

straints are asymmetric so that when house prices rise they gain equity which can be extracted, but when house prices fall to a point where households are underwater, they can continue to live in their house as long as they make fixed mortgage payments.

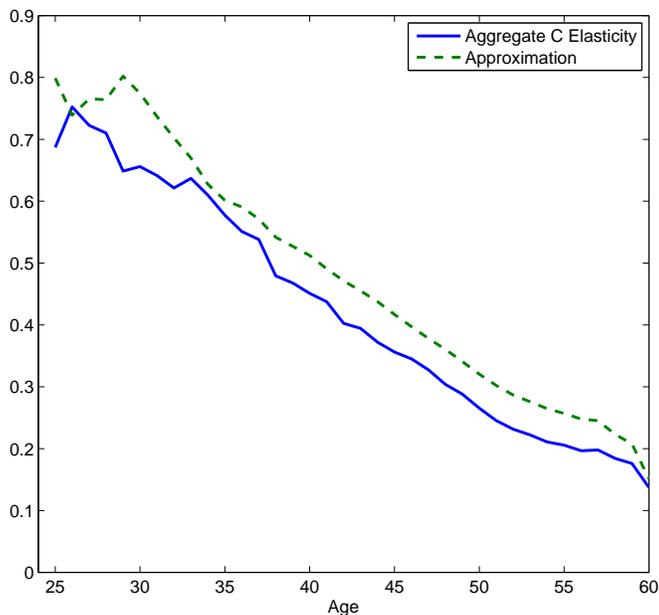
6.1 Housing Transaction Costs

To what extent does the result in Proposition 1 survive in the more general environments, which relax some of its assumptions? In this section we show that introducing realistic levels of illiquid housing, costly refinancing, an option to rent and one-sided collateral constraints do not overturn the broad conclusions developed in the frictionless environment. Once we introduce these complications, the analytical result in Proposition 1 no longer holds exactly, but we show that it remains a very good approximation for realistic calibrations of the richer model. Furthermore, implied elasticities are barely altered by the presence of illiquid housing or asymmetric collateral constraints. The option to rent has a more significant effect on the aggregate response of consumption to house price movements but implied elasticities remain substantially larger than the PIH model and are in line with empirical estimates.

We begin by relaxing the assumption that housing can be adjusted costlessly. In particular, the transaction cost of housing adjustment is given by $\kappa_{it} = FP_t H_{it-1} \mathbb{1}_{H_{it} \neq H_{it-1}}$ where $\mathbb{1}$ is an indicator function equal to 1 iff $H_{it} \neq H_{it-1}$. In the frictionless model we set $F = 0$, but we now set $F = 0.05$. Typical brokers fees in the U.S. are 6%. Splitting these evenly between the buyer and seller would imply a 3% transaction cost. However, there are typically additional closing costs together with search costs incurred when adjusting housing, so we choose a somewhat larger value. Our benchmark 5% transaction cost is exactly equal to the value of housing adjustment costs calibrated in Diaz and Luengo-Prado (2010). Berger and Vavra (2015) estimate adjustment costs of .0525 for a broad measure of durable spending. Redoing calculations with adjustment costs between 2.5% and 7.5% has no substantive effect on our conclusions.

The presence of fixed adjustment costs substantially complicates the household optimization problem as households' policy functions are now highly non-linear and characterized by (S,s) style inaction regions. This in turn implies that the previous optimality conditions and Proposition 1 no longer hold exactly. See the compu-

Figure 5: Elasticities over the Life Cycle: Baseline Model with No Rental Market

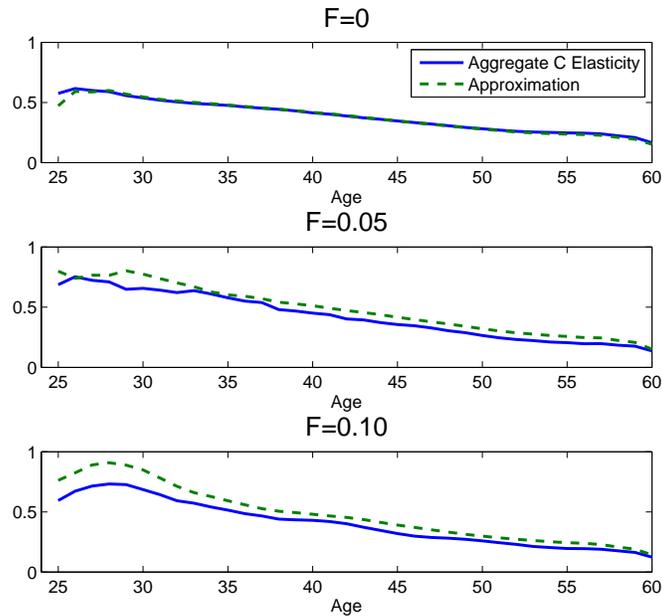


tational appendix for a detailed description of the numerical solution method. We again initialize age 25 households in the model to match the distribution of housing, liquid assets net of mortgage debt and income in the 2001 SCF and target the life-cycle profile of these variables. Overall the fit remains quite good.

What are the implications of housing transaction costs for the accuracy of our sufficient statistic and for aggregate elasticities. The blue line in Figure 5 shows the average elasticity of non-durable consumption as a function of age. This elasticity is extremely similar to that in Figure 2 for the frictionless model. The green dashed line in Figure 5 instead computes the elasticity using the sufficient statistic from Proposition 1. Since Proposition 1 no longer holds exactly, we henceforth refer to the elasticity computed using $MPC \times H$ as the "Approximation Formula". Figure 5 shows that despite the presence of illiquid housing, the approximation remains highly accurate.

This allows us to provide intuition for the high and declining elasticity that is similar to that in the frictionless model: The reason for the high elasticities for the young agents is that they have high MPCs as well as substantial housing. As households age, they accumulate assets and their MPC falls. Households can borrow to

Figure 6: Robustness to Different Fixed Costs: $\theta=0.25$, No Rental

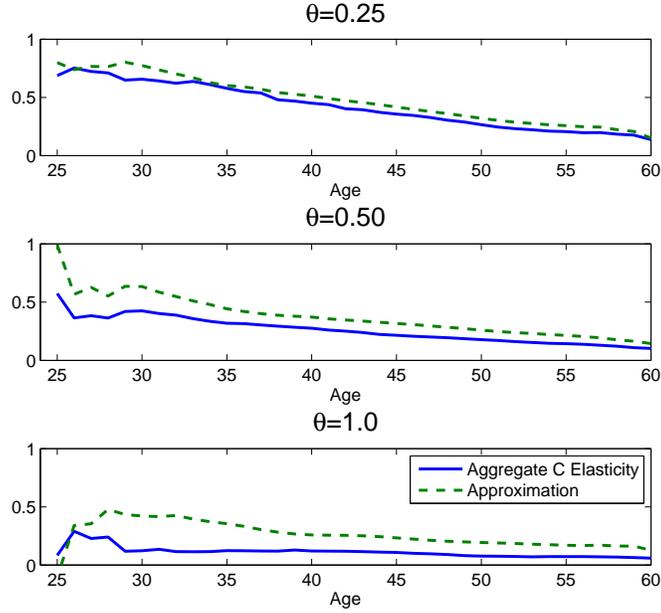


finance houses early in life, so the life-cycle profile of housing is flatter than that of the MPC. The presence of housing transaction costs further amplifies the incentive to purchase a large house all at once using leverage rather than slowly accumulating housing.

How robust is the accuracy of the approximation formula in Proposition 1 to alternative environments? Figure 6 shows results for a range of fixed costs, holding θ constant at its baseline value. Figure 7 shows results of varying θ while holding fixed costs constant at their benchmark value of 0.05.

These figures show that even for fixed costs which are substantially larger than those in the data, the approximation remains highly accurate. The accuracy of the approximation is somewhat more sensitive to the value of the downpayment, but remains highly accurate for empirically realistic values. Only when downpayments rise to 100% does the formula break down substantially. Unsurprisingly, as fixed costs are increased and it becomes harder to extract housing wealth, household behavior moves further away from the frictionless model. These results demonstrate that the result in Proposition 1 is not tautological and that it does not hold in all models. Nevertheless, it is extremely accurate for empirically realistic specifications

Figure 7: Robustness to Different Downpayments: $F=0.05$, No Rental



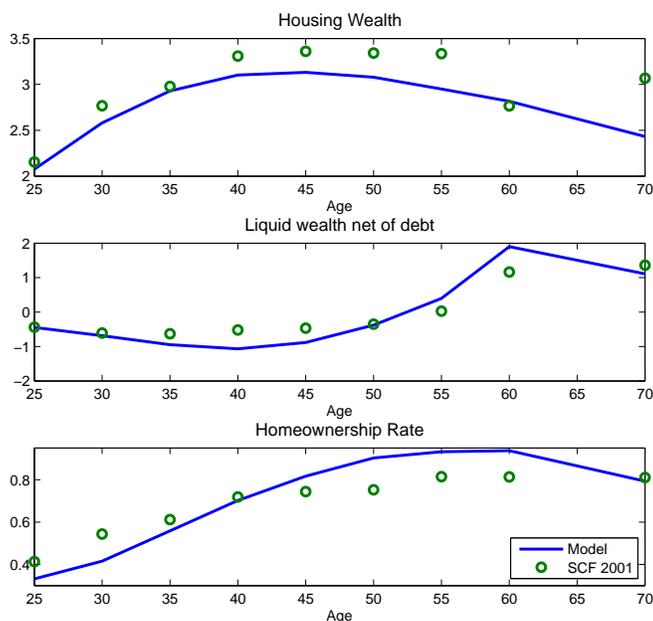
for fixed costs.

6.2 Rental Markets

In all results thus far, we have assumed that households must purchase housing in order to consume housing services. When housing can be adjusted frictionlessly, this is not a strong assumption, but once housing is illiquid this assumption is no longer innocuous: households who are liquidity constrained are essentially forced to buy households when they would prefer to rent if given the option. In this section, we introduce rental markets into the model in the previous subsection. In particular, we assume that households can rent H units of housing by paying $r_{rental} \times H$. Rental housing must be repurchased each period and cannot be used as collateral. However, rental housing can be adjusted costlessly. The flow cost of renting, r_{rental} , relative to owner-occupied housing is now an additional parameter in the model and we follow a calibration strategy for our parameters similar to that in our previous results.⁷ As before, we target the life-cycle profile of housing and liquid assets

⁷At retirement, household risk falls to zero. This substantially changes the trade-off between liquid and illiquid assets. If the rental rate of housing remains constant then there is a large jump up

Figure 8: Life-Cycle Moments: Data vs Model



net of debt, but we now add the life-cycle profile of homeownership rates as an additional moment.

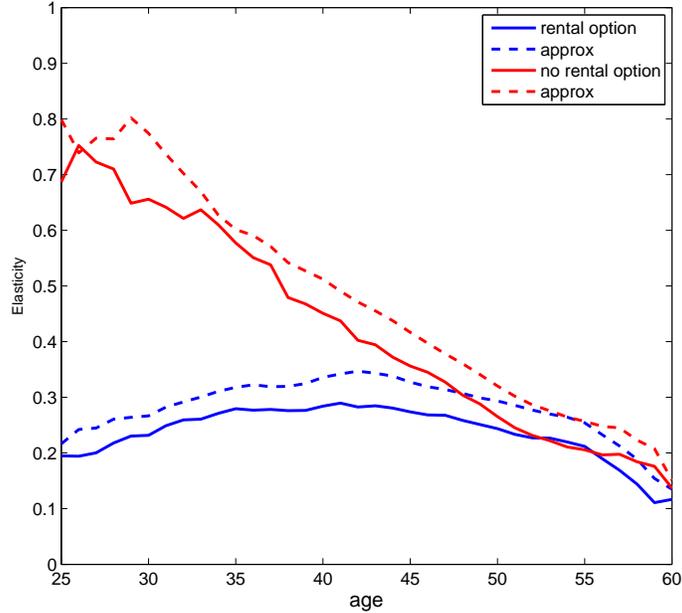
Figure 8 shows that, as before, we are able to match the life-cycle profile of housing and liquid wealth reasonably well. In addition, we can match the upward slope and subsequent flattening of homeownership rates over the life-cycle.

How does the addition of a rental option affect the response of consumption to house prices as well as the accuracy of our approximation? Figure 9 compares the true model elasticity to the approximation for the model with and without rent. There are two important takeaways from this figure. First, the addition of a rental option substantially lowers the response of consumption to house prices. However, the elasticity remains very substantial. In the model with rental, the average elasticity over working life is 0.24, as compared to a value of 0.43 for the comparable model without the option to rent.⁸ The value implied by the model with rental is

in the homeownership rate at retirement. To eliminate this jump, we introduce a different rental rate at retirement r_{rental}^{retire} . This is isomorphic to a lower utility of housing in retirement, perhaps due to no longer having children at home or to greater challenges to home maintenance. However, since we concentrate on results for working age households, this retirement parameter has little effect on any of our results.

⁸If we include retired households then the average elasticities are 0.21 and 0.40.

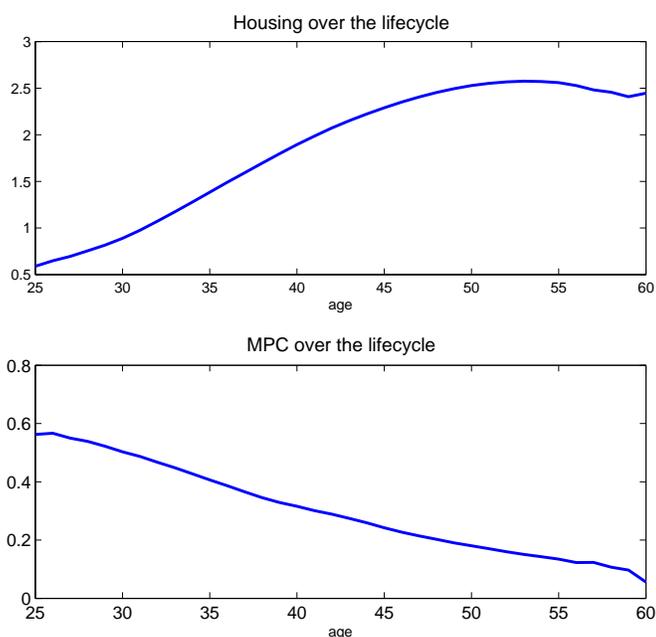
Figure 9: Elasticities over the Life Cycle



both large and completely in line with estimates from the empirical literature discussed in Section 2. If anything, the elasticity implied by the model with no rental option is too large.

The second important takeaway from Figure 9 is that the addition of a rental option does not affect the validity of the approximation. It remains the case that the approximation based on Proposition 1 remains highly accurate. This in turn allows us to immediately get intuition for why the elasticity of consumption to house prices falls when households have the option to rent. First, note that since renters have zero housing, they do not contribute to the aggregate response of consumption to house price movements $MPC \times H$. Since the average homeownership rate in the data is 70%, matching this number mechanically lowers the elasticity of consumption to house prices even if all households are identical. However, not all households are identical and homeownership is not randomly distributed in the population. Figure 10 shows the life-cycle profile of housing and MPCs in the model with the option to rent. While the young in the model have high MPCs, they have substantially less housing. In both the model and the data, the life-cycle profile of homeownership is steeply upward sloping. This substantially reduces the consumption response of the

Figure 10: Understanding the Life Cycle

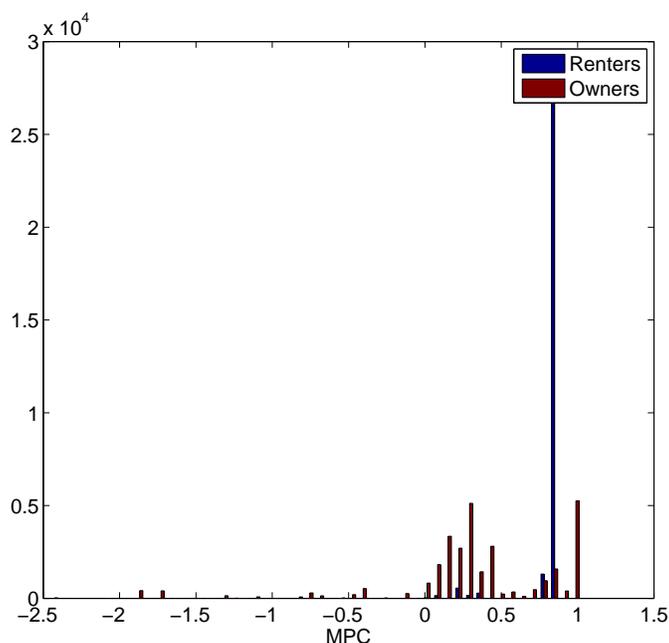


young to house price movements. Indeed, by age 45 the models with and without rental markets are nearly identical.

We can get another sense of the importance of households with high MPCs endogenously choosing to rent by comparing the distribution of MPCs for renters and owners in the model. Figure 11 shows that the MPCs for households that choose to rent are substantially higher than for households that choose to own. This is because households with high MPCs are relatively constrained and so more highly value having liquid wealth available rather than illiquid housing. Again, the approximation formula then immediately explains why this reduces consumption responses to house price movements.

How robust is the conclusion the accuracy of the sufficient statistic once we introduce the option to rent? Just as before, Figures 12 and 13 show that the formula remains accurate for a range of parameter values. The accuracy of the formula is little affected by changes in the size of housing adjustment costs. It is again somewhat more sensitive to changes in the amount of housing which can be collateralized, but is highly accurate for reasonable values of this parameter. It is also worth noting that as θ rises, the elasticity falls. This is the combination of two factors. First,

Figure 11: MPC Distribution: Renters vs Owners



household leverage falls which reduces MPCs for a given housing value. Second, in these comparative statics exercises we do not recalibrate the model to hit data moments as we vary F and θ . As we raise the required downpayment while keeping r_{rental} constant, homeownership becomes less and less attractive. In our baseline results with $F = 0.05$ and $\theta = 0.25$, we match the average homeownership rate in the data of 70%, but as we raise θ to 1.0, the homeownership rate in the model falls to 0.23. Thus, most of the decline in the elasticity as θ increases occurs because the homeownership rate is driven to counterfactually low levels.

6.3 More realistic mortgages

In this section we discuss the extension of our model to a somewhat more realistic mortgage process. While we do not attempt to model the details of 30-rate mortgages and household default, we can easily extend the model's realism in two dimensions. First, in our results thus far, we have assumed that as long as households meet the collateral constraint they can always borrow more without facing any cost. In reality, households with positive housing equity can extract it by refinancing, but there is some cost to doing so. This means that even with transaction

Figure 12: Robustness to Different Fixed Costs: $\theta=0.25$, Rental

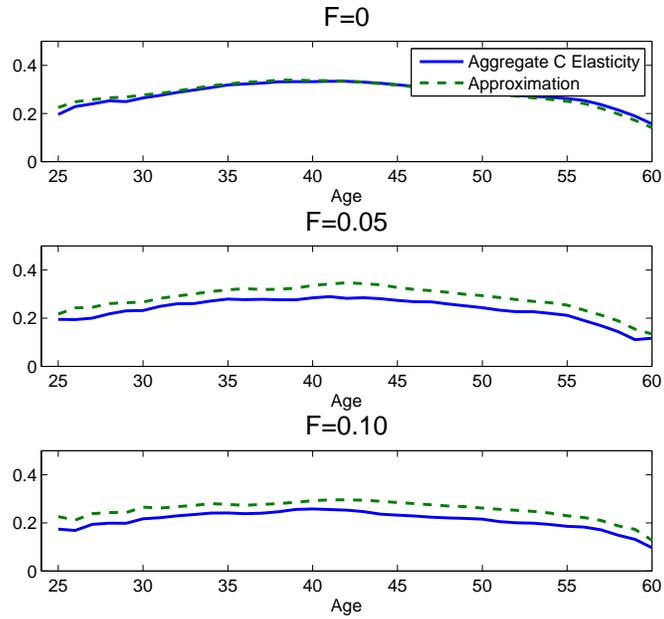


Figure 13: Robustness to Different Downpayments: $F=0.05$, Rental

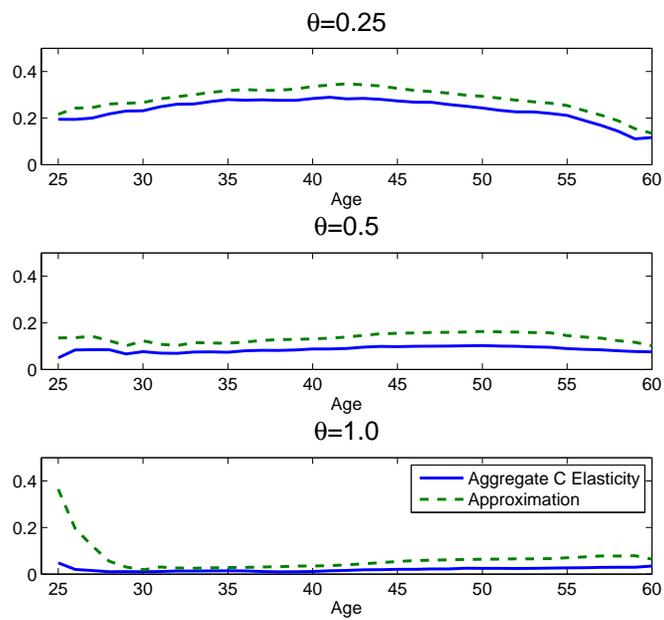
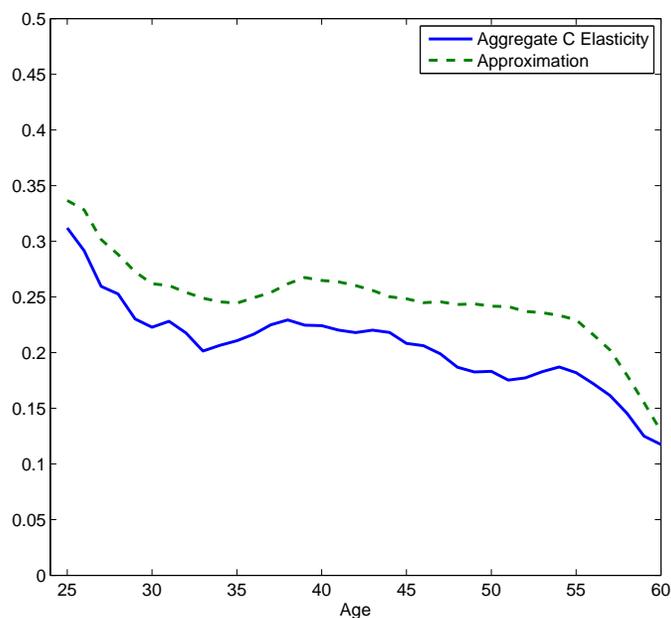


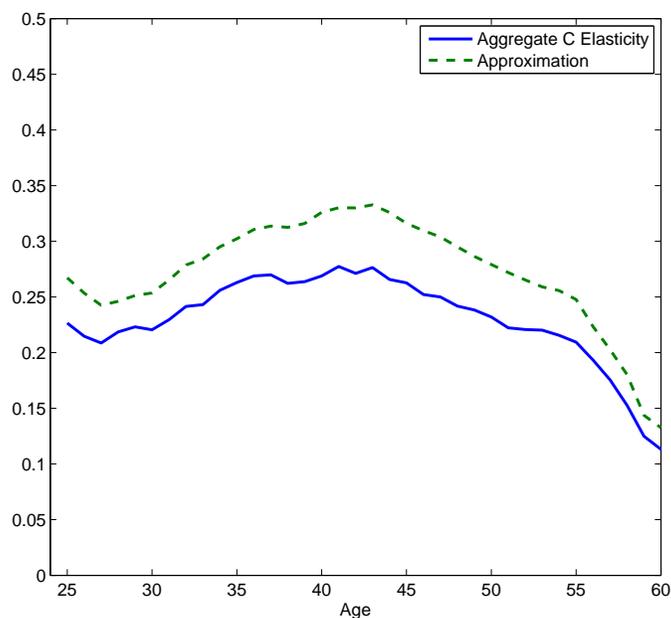
Figure 14: Elasticities over the Life Cycle: Costly Refinancing Model



costs on housing adjustment, our results thus far probably overstate the liquidity of housing wealth. In this section we extend the model so that households can always save and dis-save with no cost, but if they want to start borrowing or increase the level of their borrowing then they must pay a transaction cost to do so. We pick this cost of refinancing to match the fraction of refinancing observed in Bhutta and Keys (2014), which implies a fixed cost of 0.005. In the model with no refinancing cost, roughly 30% of households increase debt each year while with a refinancing cost of 0.005 this number falls to a realistic 10%. [results to be added] The results show that the introduction of costly refinancing does not change any of our conclusions. The sufficient statistic remains highly accurate and elasticities are barely affected. Thus, the effect of making collateralized borrowing as costly as it is in the data is fairly negligible.

In addition, mortgage contracts are asymmetric. When house prices rise and households gain housing equity, they have the option to extract this equity if they so choose. In contrast, when house prices fall, households are not required to put up additional collateral unless they want to refinance. That is, households are not required to reduce consumption when house prices fall even if this makes them vi-

Figure 15: Elasticities over the Life Cycle: Asymmetric Mortgage Contracts



olate the collateral constraint on new loans. Households always have the option of continuing their current mortgage payments and living in a home which is underwater rather than increasing equity in their house. It is again straightforward to introduce mortgage contracts of this form into the model. This again has almost no effect on our conclusions.

7 Empirics

In previous sections we have shown that our sufficient statistic continues to hold well even in realistic model environments. In this section, we attempt to measure the sufficient statistic empirically. Recall, that this requires data on both a) the MPC out of a transitory income shock and b) data on nominal home values. While it is easy to obtain information on the latter, estimating the former is much more difficult. However, once we have both pieces of information it is easy to compute the implied consumption elasticities (using our sufficient statistic), which can we use to compare to previous empirical estimates as well as to our model results.

BPP approach

We estimate the MPC out of transitory income shocks by following the identification approach of Blundell, Pistaferri and Preston (2008, henceforth BPP).⁹ BPP showed that if income follows a process with a permanent and i.i.d. component then one can identify the MPC out of the transitory shock if one has individual level panel data on income and consumption. In particular, BPP assume that log income is equal to $y_{it} = z_{it} + \varepsilon_{it}$, where z_{it} follows a random walk with innovation η_{it} and ε_{it} is an i.i.d. shock. It follows that the change in log income is equal to $\Delta y_{it} = \eta_{it} + \Delta \varepsilon_{it}$. This is a very common income process in the empirical labor literature and Abowd & Card (1989), who showed that this specification is parsimonious and yet fits income data well.

BPP showed that given this income process, the true MPC out of transitory shock is equal to:

$$MPC_t = \frac{cov(\Delta c_{it}, \varepsilon_{it})}{var(\varepsilon_{it})}$$

More importantly, under the assumptions that the household has no foresight or no advanced information about future shocks, a consistent estimator of the MPC is:

$$\widehat{MPC}_t = \frac{cov(\Delta c_{it}, \Delta y_{it+t})}{cov(\Delta y_{it}, \Delta y_{it+1})}$$

If one has access to panel data with at least three time periods one can easily implement this estimator by running an instrumental variable regression of the change in consumption Δc_{it} on the change in income Δy_{it} , instrumenting for the current change in income with the future change in income, Δy_{it+1} . Since we need individual level panel data on income and consumption, we use data from the Panel Study of Income Dynamics (PSID).

An obvious drawback of using the BPP approach is that one must take a stance on the exact parametric form of the income process. An alternative approach would be to estimate the MPC to a transitory income shock in the Consumer Expenditure Survey (CEX) using randomization in the timing of the receipt of government tax rebates, an approach that has been popularized by Johnson, Parker and Souleles (2006), henceforth JPS). There are a number of reasons why we prefer using the BPP approach. First, the PSID has significantly larger sample sizes (35K versus 10K)

⁹For more details about how this identification scheme please see Kaplan and Violante (2009) and Kaplan, Violante and Weidner (2014).

than the CEX. More importantly, the panel structure of the PSID allows us to more easily tell whether large changes in consumption reflect true changes in the desired consumption of individuals rather than measurement error. Finally, the coverage of consumption categories in the PSID is both more complete than in the CEX (mean consumption to income ratios are respectively 49% and 35%), but the coverage in the PSID is also a constant fraction¹⁰ of NIPA consumption. This is important because it is well known that ratio of CEX consumption to NIPA measures of consumption has been declining over time, most likely due to the difficulty in measuring the consumption of the rich. Nonetheless, despite these caveats we have also repeated our results using JPS's data and found qualitatively similar results.

PSID data

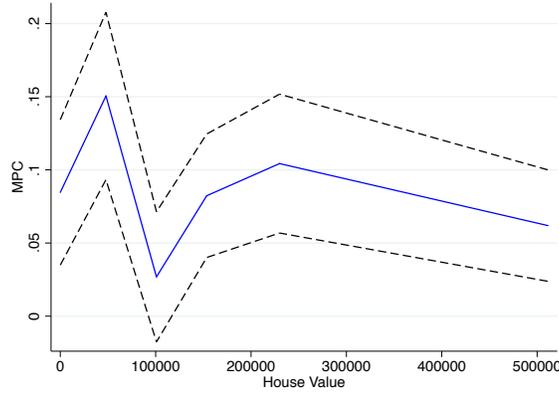
Implementing our sufficient statistic empirically requires a longitudinal dataset with information on income, consumption, and housing values at the household level. Starting from the 1999 wave, the PSID contains the necessary data. The PSID started collecting information on a sample of roughly 5,000 households in 1968. Thereafter, both the original families and their split-offs (children of the original family forming a family of their own) have been followed. The survey was annual until 1996 and became biennial starting in 1997. In 1999 the survey augmented the consumption information available to researchers so that it now covers over 70 percent of all consumption items available in the Consumer Expenditure Survey (CEX).

Since we use almost the same underlying sample as Kaplan, Violante and Weiland (2014), our description of the PSID closely mirrors theirs. We start with the PSID Core Sample and drop households with missing information on race, education, or state of residence, and those whose income grows more than 500 percent, falls by more than 80 percent, or is below \$100. We drop households who have top-coded income or consumption. We also drop households that appear in the sample fewer than three consecutive times, because identification of the coefficients of interest requires a minimum of three periods. In our baseline calculations, we keep households where the head is 25-75 years old. Our final sample has 39,722 observations over the pooled years 1999-2011 (seven sample years).

We use the same consumption definition as Blundell, Pistaferri and Saporta-Eksten (2014), which covers approximately 70% of total NIPA consumption. Our

¹⁰Around 70%. See table 2 in Blundell, Pistaferri and Saporta-Eksten (2014).

Figure 16: MPC and Housing



definition of income is defined as the sum of labor income and government transfers. We purge the data of non-model features by first regression Inc_{it} and lny_{it} on year and cohort dummies, education, race, family structure, employment, geographic variables, and interactions of year dummies with education, race, employment, and region.

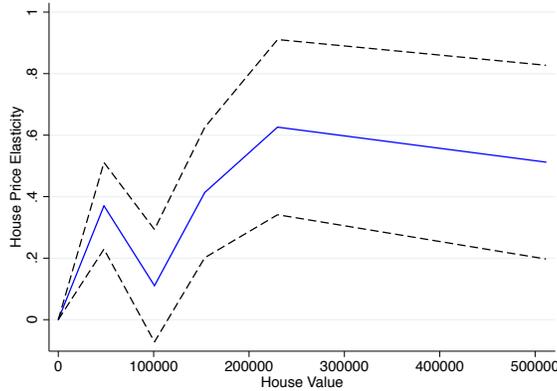
Results

Since the consumption elasticity is equal to the product of the MPC and the value of housing, we first examine how the MPC varies in response to the value of housing. To do this, we estimate the MPC using the BPP methodology separately in 6 housing bins. The first and largest bin includes only renters - that is those households who by definition own zero housing. The other five bins are constructed as the quintiles of the home value distribution and so are equally weighted by construction. In total, we have approximately 15,000 renters in our database and 25,000 home owners.

Figure 16 shows the relationship between the MPC and home values. Overall the relationship between the MPC and home values is flat or slightly declining (though this decline is not statistically significant). Figure 17 shows the relationship between the consumption elasticity and home values where the consumption elasticity is measured using the formula for our sufficient statistic. That is the elasticity of nondurable consumption to house prices is:

$$\eta_{it} = MPC_{it} \cdot \frac{(1 - \delta) P_t H_{it-1}}{C_{it}}.$$

Figure 17: Elasticity and Housing

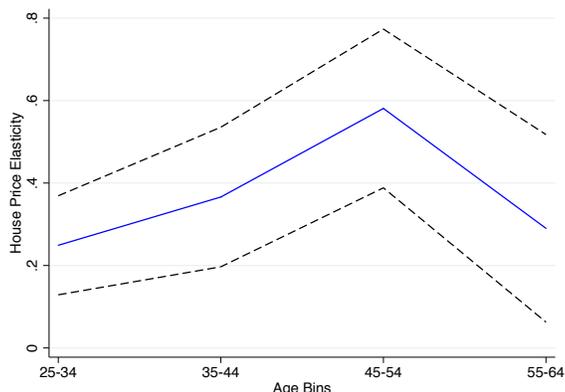


where i denotes each housing bin and $P_t H_{it-1}$ and C_{it} are the mean level of home values and nondurable consumption in each bin.

Unsurprisingly, the implied elasticity is increasing over much of the home value range since the formula depends positively on nominal home values. One interesting statistic to compute is the implied average nondurable consumption elasticity. This value is computed by multiplying the value of the elasticity in each bin by the fraction of agents in each bin. The value of this elasticity in the PSID data is 0.26, which is consistent with the range of values that Mian, Sufi and Rao (2013) find. It is interesting to compare the above value of the nondurable consumption elasticity to one finds if you compute the average elasticity using the entire data set and only use overall mean levels of nondurable consumption and home values. This calculation is interesting because tells you how important heterogeneity in homeowner status is in the measurement of the consumption elasticity. Our results using our general model suggested this heterogeneity was quite important because households with the highest MPCs were the least likely to own housing.

Indeed, we find evidence that this taking into account this heterogeneity is important. The overall elasticity one finds when abstracting from homeowner status is equal to 0.34 with a 95% of (0.26,0.43). This is 30% higher than the estimate we find when we take into account heterogeneity in homeowner status. Next, we explore the relationship between the nondurable consumption elasticity and the lifecycle. Ideally, we would like to estimate the elasticity for all ages 25-65 in order to ease comparison with our model result. Due to sample size limitations, however, we instead compute results for four age bins. These bins are all households aged 25-34,

Figure 18: Elasticity and Age

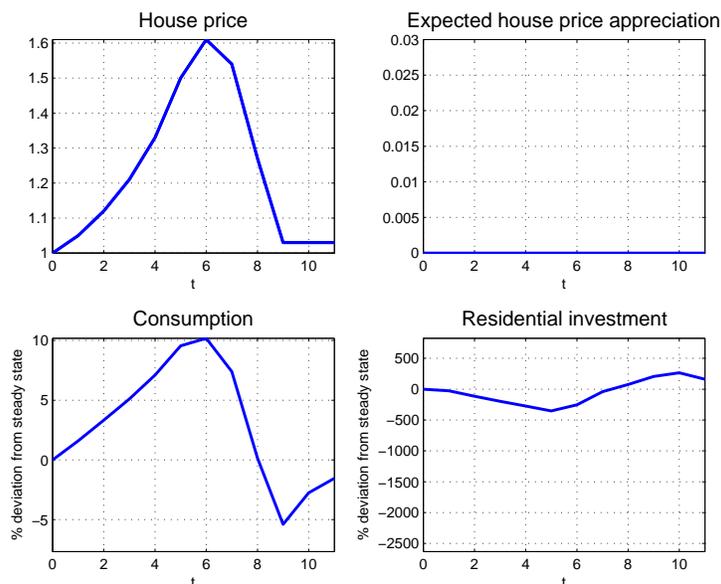


35-44, 45-54 and 55-64, respectively. The results showing the relationship between the elasticity and the lifecycle are shown in figure 18. Our results suggest that the relationship between the elasticity and the lifecycle is hump shaped. This hump shape reflects two separate forces. First, MPCs are hump-shaped in age (potentially) because many middle aged households are wealthy hand to mouth consumers in the parlance of Kaplan and Violante (2014). That is, many middle aged households hold significant illiquid wealth holdings but small amounts of liquid wealth which limits their ability to smooth income and wealth shocks. Second, homeownership rates are hump shaped so mechanically the elasticity is going to be hump shaped. Finally it is worth comparing figure 18 and figure 9, which shows the relationship between the nondurable consumption elasticity and the lifecycle in our baseline model. It is striking how similar our baseline model results are with our empirical estimates despite the fact that our baseline model still abstracts from many features of the housing market. Both figures show a prominent hump shape and the magnitude of elasticities are quite close.

8 A “Bubble” Experiment

We now approach the question whether an heterogenous agents model like the one developed in this paper can be used to shed light on the behavior of aggregate consumption in housing boom-bust episodes, like the one experienced by the US in the 2000s. A central feature of housing boom-bust cycles is that residential investment moves in the same direction as house prices. In the simple partial equilibrium

Figure 19: Bubble experiment: no expected appreciation



model developed so far, an increase in house prices generates a drop in housing demand and thus a drop in residential investment.¹¹ Therefore, the model needs to be enriched to capture the procyclical behavior of residential investment. In particular, we need some force that drives households' demand for housing up in the boom and down in the bust. Here we experiment with changes in households' expectation of future home price growth. Recent empirical work by Case, Shiller and Thomson (2012) provides survey evidence supportive of the view that there was an increase in expected house price appreciation during the recent US boom.

Our experiment is as follows. We start the economy in a steady state with constant house prices. We then simulate the model for T periods, with a given path of realized prices $\{P_t\}_{t=0}^T$. In each period t , we assume that households observe the price today P_t and expect a deterministic path of house prices in the future, given by

$$E_t[P_{t+j}] = P_t \exp\left(g_t \frac{1 - \lambda^j}{1 - \lambda}\right),$$

for $j = 1, 2, \dots$, so households expect prices to grow by $\exp g_t$ (in gross terms) in the next period and to keep growing gradually towards the long-run level $P_t \exp(\frac{g_t}{1-\lambda})$.

¹¹For some, highly leveraged households, housing demand can actually increase as a result of a house price increase. However, with any reasonable parametrization of initial conditions aggregate housing demand goes down when P increases.

Households' optimal behavior at time t gives us consumption and housing demand and determines the endogenous evolution of the distribution of housing and non-housing wealth in period $t + 1$. We dub this experiment a "bubble" simply because house price expectations are not disciplined by rational expectations. The sequence $\{P_t\}_{t=0}^T$ is chosen to match the US price series from Shiller (2015).¹² The choice of the sequence $\{g_t\}_{t=0}^T$ is discussed below.

Let us consider the model with no adjustment costs of Section 5. Consider first the case in which consumers receive a sequence of unexpected price shocks that are always perceived as permanent. That is, consider the case in which $g_t = 0$ for all t . Figure 20 shows that in this case, the response of consumption can be quite large. Notice that the effects are non-linear and state-dependent. The first shock, at date 0, has an effect which is in line with the elasticity reported in Section 5: the house price increase is 5% and the consumption increase is 1.7% = 0.34 · 5%. However, the cumulated effect of the shocks is smaller, so a 60% appreciation leads to a 10% increase in consumption. The reason is that as consumers' net worth increases, MPCs tend to decline. On the other hand, when house prices fall, the elasticity is slightly larger than 0.34.¹³ However, the bottom-right panel of the figure shows that the model predictions for housing demand are clearly counterfactual, with residential investment falling in the boom and increasing in the bust.

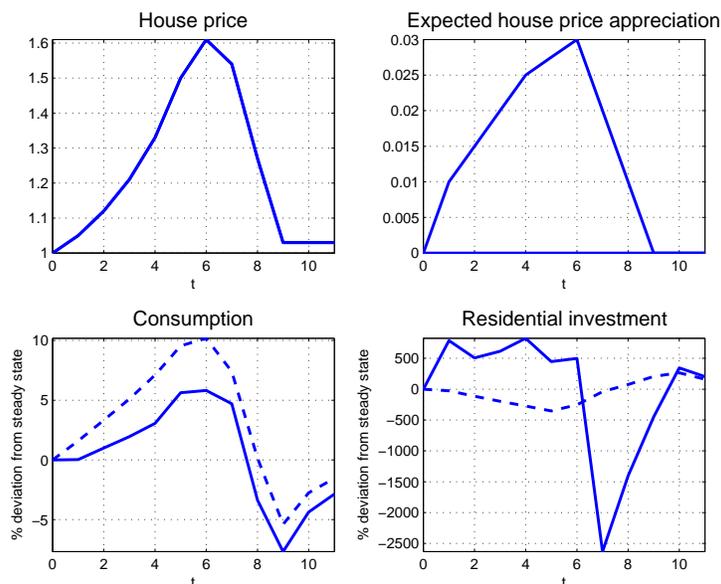
We now introduce changes in expected appreciation and explore what are the consequences for our predictions on consumption. In Figure ?? we report the results of an experiment in which we feed the model a sequence g_t that displays mild appreciation expectations building up in the boom phase, which then go away gradually in the bust. The dashed lines in the figure replicate the effects of our first experiment with $g_t = 0$, for comparison. The effects on housing demand are very large and now we obtain fluctuations in residential investment that go in the right direction qualitatively. Quantitatively the effects are clearly too large. Since housing is perfectly liquid in this version of the model, housing demand is very sensitive to changes in the user cost of housing, which, in turn, is very sensitive to changes in g_t .

Turning to the model predictions for consumption, now the effects are smaller during the boom and of similar magnitude in the downturn. Several forces explain this different behavior. On the one hand, a change in g_t changes consumption poli-

¹²We use the Real Home Price Index from www.econ.yale.edu/shiller/data, file Fig2-1.xls.

¹³The consumption percentage change is .15/1.1 and the house price percentage change is .6/1.6, so the elasticity is 0.36.

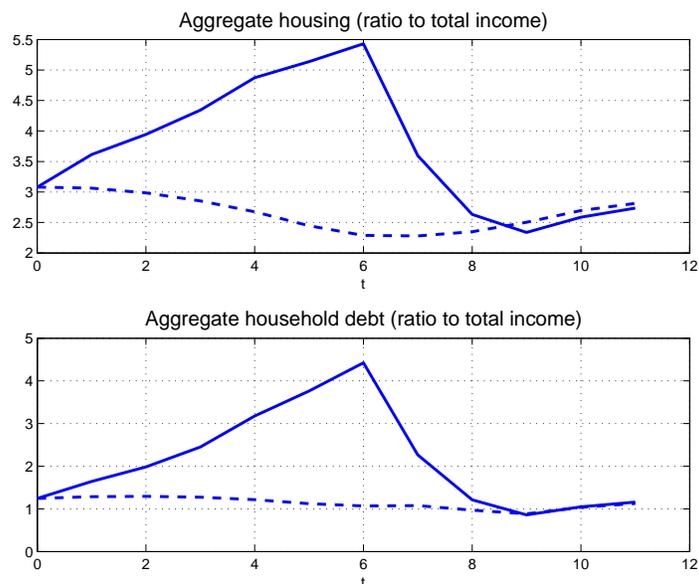
Figure 20: Bubble experiment: with expected appreciation



cies. In particular, in our model an increase in g_t tends to reduce consumption for given W , when $\sigma > 1$. The reasoning is as follows: an increase in g_t implies an expected increase in the cost of housing services in the future and thus an expected increase in the overall CPI. This has an intertemporal substitution effect which tends to increase consumption today (when the CPI is lower) and an income effect, which tends to lower consumption today. When $\sigma > 1$ the income effect dominates and consumption falls. Summing up, in the boom phase we now have two forces: the “wealth effect” explored in previous sections, which tends to increase consumption, and the effect of g_t , which tends to decrease consumption. The wealth effect is sufficiently large that overall we still obtain a fairly large consumption increase.

The same forces are at play in the downturn. However, in the downturn the presence of g_t now also has an amplifying effect. While the effect of g_t on the consumption policies is still to dampen the effects (now a reduction in g_t stimulates consumption today, thus dampening the fall), we now have to take into account the effect of the g_t 's in the boom phase on the household balance sheets when they enter the downturn. Coming from a period of high g_t 's means that households enter the downturn with a much higher level of housing and debt, so a negative shock to housing values has a bigger proportional effect on their net worth. This amplifying effect working through the state variables roughly balances the dampening

Figure 21: Bubble experiment: housing and debt



effect working through the consumption policies, thus explaining why the fall in consumption is similar in the two exercises. This interpretation is supported by Figure 21 which reports the behavior of the housing stock and of household debt in the two experiments.

The preliminary take-out from these exercises is as follows. Introducing shocks to expected appreciation can deliver a boom-bust in residential investment together with a boom-bust in house prices. Shocks to expected appreciations change the consumers' intertemporal problem, so when we look at consumption responses (in the model and in the data) we have to interpret them as the *joint* responses to house price level changes and to changes in expected future prices. On the one hand, shocks to expected appreciation tend to dampen the boom-bust in non-durable consumption through intertemporal channels (with a low elasticity of intertemporal substitution). On the other hand, they tend to amplify the boom-bust through endogenous leverage decisions. It remains to analyze how these forces play out in the full model with adjustment costs and rentals.