

House Prices and Consumer Spending *

David Berger

Northwestern University

Veronica Guerrieri

University of Chicago

Guido Lorenzoni

Northwestern University

Joseph Vavra

University of Chicago

November 21, 2016

Abstract

Recent empirical work shows large consumption responses to house price movements. Can consumption theory explain these responses? We consider a variety of consumption models with uninsurable income risk and show that consumption responses to permanent house price shocks can be approximated by a simple “sufficient-statistic” formula: the marginal propensity to consume out of temporary income times the value of housing. Calibrated versions of the models generate house price effects that are both large and sensitive to the level of household debt in the economy. We apply our formula to micro data to provide a new measure of house price effects.

Keywords: Consumption, House Prices, MPC, Leverage, Debt, State-Dependence.

JEL Codes: E21, E32, E6, D14, D91, R21.

*Email: Berger: david.berger@northwestern.edu; Guerrieri: vguerrie@chicagobooth.edu; Lorenzoni: guido.lorenzoni@northwestern.edu; Vavra: joseph.vavra@chicagobooth.edu. We thank Sasha Indarte and David Argente for excellent research assistance. We would also like to thank Orazio Attanasio, Adrien Auclert, João Cocco, Jonathan Parker, Monika Piazzesi, Martin Schneider, Alp Simsek, Chris Tonetti, Gianluca Violante and seminar participants at Harvard, MIT, Rochester, Duke, CREI, Johns Hopkins, Brown, Cambridge CFM, CSEF-CIM-UCL Workshop, Empirical Macro Workshop-Austin, Maryland, Northwestern, Yale, UCLA, Penn State, SED-Warsaw, EIEF, SITE, Stanford, LSE, LBS, Bocconi, Boston Fed, St. Louis Fed, Kansas City Fed, Boston University, Michigan, Berkeley and Booth.

1 Introduction

The last US recession was characterized by a large and persistent decline in consumer spending.¹ A growing empirical literature has argued for the importance of house prices and household debt in explaining consumption patterns both before and during the recession. More broadly, there is widespread policy concern that boom-bust cycles in housing and consumer debt can end with large contractions in consumer spending.² However, the theoretical rationale for house price effects on consumption is less clear. In particular, it is a commonly held view that these effects should be small, because increases in the value of an individual's house are offset by increases in future implicit rental costs, leaving the lifetime budget constraint unchanged. If households make consumption decisions based on the lifetime budget constraint, as in a permanent income hypothesis (PIH) model, then consumption effects are small.³

In this paper, we explore whether models can deliver the large effects found in the empirical literature and what features are important for doing so. Are house price effects larger when the household sector is more levered? Is it important that housing can be used as collateral? Can house price movements independently lead to large consumption swings or do they require correlated movements in common factors such as income or financial conditions?

We study these questions using an intentionally standard class of incomplete market models with income uncertainty and housing that serves as collateral which includes PIH and many other classic models as special cases. Our main quantitative conclusions are: (i) In contrast to PIH intuition, realistic model calibrations produce large consumption responses to house price changes, in line with the recent empirical literature. (ii) The size of these responses depends crucially on the economy's joint distribution of housing and debt.

What generates large consumption responses to house price movements? In our framework, increases in house prices affect consumption through four channels: 1) A positive "endowment" effect which arises from the increase in the value of an individual's current house. 2) A negative "income" effect which arises from the increase in current and future implicit rents. 3) A positive "substitution" effect as households substitute from housing to consumption. 4) A positive "collateral" effect as increases in house prices relax households' borrowing constraints.

Given these competing effects, one might expect that consumption responses would be quite complicated to characterize, but we show this is not the case. The central theoretical result of our paper is that individual consumption responses to unexpected, permanent house price changes are given by a simple "sufficient-statistic" formula: the individual marginal propensity to consume

¹See Petev et al. (2012).

²See Cerutti et al. (2015).

³Sinai and Souleles (2005) p. 773 clearly formulate this view: "increases in house prices reflect a commensurate increase in the present value of expected future rents" and "for homeowners with infinite horizons, this increase in implicit liabilities would exactly offset the increase in the house value, leaving their effective expected net worth unchanged." Similar arguments are made in Glaeser (2000), Campbell and Cocco (2007) and Buiter (2008).

out of temporary income shocks (MPC) times individual housing values. The aggregate response is then determined by the endogenous joint distribution of MPCs and housing. Since our modeling environment is intentionally standard and includes many simple consumption models as special cases, this means that our sufficient-statistic can be used to explain exactly why different models and assumptions lead to very different implications for consumption responses to house price changes. For example, PIH models with no borrowing constraints generate small MPCs which have little correlation with housing. In models with borrowing constraints, MPCs increase, and the presence of leverage makes MPCs highly correlated with housing values. Our formula also explains why, following different shock histories which generate different distributions of housing and debt, an economy displays different aggregate responses to current shocks.

Our sufficient-statistic tells us that the consumption response to a house price change is determined by its effect on household endowments times the marginal propensity to consume. This is the definition of the endowment effect, and so it follows that we can capture the total house price effect on consumption through its pure dollar effect on the current budget constraint. While it is not surprising that house price changes deliver an endowment effect through the *current* budget constraint that takes the same form as a dollar transfer,⁴ the surprising insight is that *future* budget constraint effects, substitution effects and collateral effects exactly cancel.⁵ Thus, despite their potential complexity, one can interpret total house price effects on consumption just like pure dollar transfers.

Our decomposition also provides additional intuition for why models with borrowing constraints imply much larger consumption responses than PIH models. With borrowing constraints, the timing of wealth effects rather than just their present value becomes relevant. A permanent increase in house prices immediately increases current resources, while higher implicit rental costs occur in the future. Unlike in a PIH model, borrowing constraints imply that the current resource effect is dominant for current consumption choices.⁶

Our baseline is a partial equilibrium model calibrated to life-cycle wealth data from the Survey of Consumer Finances. We show that this model delivers an average elasticity of consumption to house price shocks of around 0.5. This elasticity is large, even relative to the empirical evidence outlined below that typically estimates elasticities of around 0.2, and especially relative to PIH models that generate elasticities near zero.

Two model features explain this large elasticity. First, income uncertainty and precautionary savings mean that low net worth agents have large MPCs. Second, housing services enter the

⁴It is widely recognized that the theoretical MPC to liquid wealth shocks is the same as the MPC to temporary income shocks, and housing wealth is liquid in our baseline model.

⁵This does not imply that these other effects are small in isolation, just that they cancel out. This is important for interpreting various empirical and theoretical results in the literature.

⁶This holds even with a constant borrowing constraint, but the fact that houses serve as collateral generates an additional consumption effect, as increases in house prices relax borrowing constraints. In fact, Sinai and Souleles (2005) were careful to acknowledge the potential role of borrowing constraints and substitution effects in undermining their result.

utility function, which implies that housing wealth is not proportional to total net worth. In particular, low net worth agents borrow and so hold housing that is a multiple of net worth. Therefore, the model generates a joint distribution with a substantial fraction of households with large housing values, low net worth, and high MPCs. Using our formula, this implies that the average consumption response to house price shocks is large. Thus, standard incomplete markets models imply large causal responses of consumption to house price changes and do not require additional features such as house price movements which are correlated with future income or other fundamentals to generate these effects.

Our baseline model has no housing transaction costs, long-term debt or rental option. Therefore, it produces an unrealistic frequency of housing transactions and cannot match the large fraction of renters in the population. We extend the model to allow for illiquid housing, long-term mortgage debt with costly refinancing and rentals and show that our sufficient-statistic no longer holds exactly. However, for realistic parameter choices it still provides a very good approximation. Somewhat surprisingly, the presence of housing adjustment costs and more realistic mortgages has only minor implications for the effect of house prices on consumption. The main quantitative difference between the baseline and extended model comes from the presence of renters. In the model with renters the average response over working life falls to around 0.25. This decline occurs both because some agents do not own a house, and thus have zero endowment effect, and because agents who choose to rent are precisely the low net worth, high-MPC agents who would have the biggest responses if they owned.

While the primary contribution of our paper is not empirical, it does have implications on this front. After showing that the sufficient-statistic approximation is robust to various model specifications, we take seriously the observation that its components are observable in the data. Estimating the distribution of $MPC \times PH$ in the data (where PH denotes house values) provides an alternative sufficient-statistic-based measure of housing price effects and can also be interpreted as a test of our structural consumption model. The main challenge is to estimate MPCs conditional on housing holdings, which we address by extending the approach of Blundell, Pistaferri and Preston (2008) to recent PSID data. We find that the micro patterns of $MPC \times PH$ are similar to those implied by our simulations and that our sufficient-statistic-based measure of housing price effects is in line with existing estimates using completely different identification strategies. Thus, our results complement and confirm those delivered by alternative methods.⁷ Interestingly, we also find that splitting the PSID sample into the housing boom and bust delivers house price effects which are substantially larger in the bust.

In the final section of the paper, we show that this arises naturally in our model by extending the analysis beyond one-time partial equilibrium house price shocks to a general equilibrium exploration of boom-bust housing cycles. A typical feature of housing cycles is that residential in-

⁷As we discuss in the body of the paper, our approach has both advantages and disadvantages relative to these alternatives, so we view it as a complement rather than a substitute for this evidence.

vestment first increases and then decreases, along with house prices, suggesting an important role for shifts in housing demand. In order to generate such shifts, we introduce changes in expected *future* house price growth in order to generate equilibrium changes in residential investment which match the data. While these shifts in expectation are crucial for matching residential investment, we show that they have only mild direct effects on consumption: consumption is primarily driven by the partial equilibrium house price level effects investigated in the rest of the paper.⁸ However, this does not imply that shifts in housing demand are irrelevant for consumption dynamics. The important channel here is that past shifts in housing demand affect the joint distribution of housing and debt, and thus determine the magnitude of the house price effect. In particular, in our simulations, the increase in housing demand during the boom leads to increased leverage, and this contributes substantially to the consumption contraction caused in the bust.

The boom-bust simulations show that the distribution of housing and debt matters for the strength of house price effects. This observation is important when interpreting empirical evidence from different time periods and when contemplating potential policy intervention into housing markets. Both shocks and policy interventions may have effects on consumption that differ dramatically depending on the distribution of household state-variables at the time the policy is enacted. In this sense, our result joins a growing literature arguing that the economy may exhibit time-varying responses to aggregate shocks.⁹

Our paper is motivated by the empirical literature studying house price effects on consumption. While methodological and data differences have led to a wide range of estimates for the relationship between house prices and consumption, the literature has generally found strong relationships. Whether or not such relationships are causal is more contentious, but recent papers with new sources of identification have argued for such causality.¹⁰ We contribute to this debate both by showing that theory is consistent with large causal effects and by constructing a new sufficient-statistic based measure of these housing price effects.

The first empirical studies of the relationship between consumption and house prices focused on aggregate data and typically found elasticities of 0.1 to 0.2.¹¹ The greatest challenge for these studies is finding exogenous variation in house prices which can be used to separate direct house price effects from the effects of common factors. In particular, an important concern is

⁸These results also show that our conclusions are robust to relaxing the assumption of permanent shocks.

⁹See e.g. Berger and Vavra (2015) for applications to durable goods, Vavra (2014) for applications to prices and Caballero and Engel (1999) and Winberry (2015) for applications to investment.

¹⁰While there are different ways of measuring housing price effects, for consistency we will focus the discussion on estimates of the consumption elasticity to house prices—the percentage change in non-durable consumption due to a one percent change in house prices.

¹¹Case et al. (2013) find elasticities from 0.03 to 0.18, with most of their estimates centered around 0.10, while Carroll et al. (2011) find an immediate (next-quarter) elasticity of 0.047, with an eventual elasticity of 0.21. They report results in terms of MPCs of 2 and 9 cents (in 2007 dollars) respectively. In 2007, housing assets from the Flow of Funds were \$22,830.5 billion and personal consumption expenditures from the BEA were \$9,750.5 billion. Multiplying MPCs by this ratio delivers the reported elasticities.

that expectations about future income growth may drive both house prices and consumption.¹² This has led to a surge in interest in micro level evidence, which allows for more variation in consumption and house prices and for alternative sources of identification. For example, Mian et al. (2013) make progress on identification by using high quality credit card data together with the Saiz (2010) housing supply instrument to isolate the effects of exogenous changes in house prices on highly localized measures of spending.¹³ Their baseline estimates of the non-durable consumption elasticity are between 0.13 and 0.26.¹⁴ We view the elasticities from Mian et al. (2013) as the closest empirical analogue to the direct house price effects in our theory, so we will often compare our theoretical results to these numbers. Several recent papers such as Kaplan et al. (2015) and Stroebel and Vavra (2015) also arrive at similar numbers using new scanner spending data and additional identification strategies.

Our paper is part of a large and growing literature studying the theoretical response of consumption to house prices in quantitative models with heterogeneous agents. A number of papers have calibrated and simulated increasingly rich models with housing and debt, using them to both address aggregate questions and draw cross sectional predictions, e.g., Carroll and Dunn (1998), Campbell and Cocco (2007), and Attanasio et al. (2011). More recently, models with these features have been embedded into general equilibrium frameworks, to study the role of households' balance sheets and debt capacity in the Great Recession.¹⁵ In particular, several papers have pointed to house values as prime determinants of households' debt capacity.¹⁶ Gorea and Midrigan (2015) studies the effects of illiquid housing on consumption and savings. Huo and Ríos-Rull (2013) and Kaplan et al. (2015) are two recent papers that study heterogeneous agent general equilibrium models with endogenous house prices, consumption and income.

We view our theoretical analysis as highly complementary to this line of work. Our sufficient-statistic formula and decompositions help identify the channels at play and show the crucial role of the endogenous distribution of housing and debt for the size of house price effects. The majority of our analysis is conducted in partial equilibrium. As emphasized by Kaplan et al. (2015), house prices are equilibrium objects, which complicates the interpretation of correlations between house prices and consumption. This is because structural shocks which move house prices may themselves have direct effects on consumption so that the simple correlation between house prices

¹²See Attanasio and Weber (1994).

¹³See also Campbell and Cocco (2007) and Attanasio et al. (2009) for other recent micro studies.

¹⁴They 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. Although these estimates can be interpreted as consumption responses to exogenous house price shocks, it is important to note that they are not pure partial equilibrium responses, since they reflect both direct house price effects plus any local general equilibrium responses. In particular, they include the additional effect on consumption due to increases in local incomes driven by greater non-tradable spending.

¹⁵Good examples are Favilukis et al. (2015) and Chen et al. (2013). Early work in this direction—that does not model housing—includes Hall (2011), Guerrieri and Lorenzoni (2011), Eggertsson and Krugman (2012).

¹⁶See, Midrigan and Philippon (2011) and Justiniano et al. (2015).

and consumption will reflect both the causal effect of house prices on consumption plus any confounding effect from the underlying structural shock. In this paper, we are interested in whether housing markets themselves play an important role in shaping consumption and propagating underlying shocks. As such, we want to isolate the causal effects of house price movements, which fundamentally require modeling housing, from any confounding effects of structural shocks on consumption, which occur independently of housing. The partial equilibrium analysis used in much of our paper holds constant all possible confounding effects, and so shows that there is indeed an important causal effect of house prices on consumption in many models.

Importantly, our sufficient-statistic formula measures the strength of this causal effect and how it changes with economic conditions or across models. It does not provide a formula for the simple correlation between house prices and consumption in equilibrium, which will depend on the underlying structural shocks which drive house prices. Nevertheless, our boom-bust exercise shows that our formula can even match this simple correlation when house prices are driven by expectations of future house price growth, since these structural shocks have little direct effect on consumption. Interestingly, Kaplan et al. (2015) run a horse-race between various structural shocks and conclude that exactly this type of shock best explains the data.¹⁷

Our emphasis on sufficient statistics connects our paper to recent work by Auclert (2015). Work in public finance has widely developed the use of sufficient statistics to characterize welfare effects and optimal policy (see Chetty (2009)). In macro, the idea is to use this approach to express some aggregate response, which may be hard to measure, in terms of individual responses, which may be easier to identify in micro data. Of course, some further steps may be necessary to translate these aggregate partial equilibrium responses into general equilibrium effects. However, we see this as a promising avenue to investigate increasingly complex heterogeneous agents models and connect them to the data.

The remainder of the paper proceeds as follows: In Section 2 we present the baseline model, derive our sufficient-statistic formula, and show baseline results. In Section 3 we show that our sufficient statistic remains accurate for the extended model that includes adjustment costs, renters and more realistic mortgages and discuss the effects of these features on the size of housing price effects. In Section 4 we take our formula directly to micro data to construct a new empirical measure of house price effects. Finally, in Section 5 we simulate a housing boom-bust episode in equilibrium and show that it leads to substantial time-variation in the strength of house price effects.

¹⁷They also show that our sufficient-statistic formula can match the correlation between house prices and consumption in their model in response to this shock, again since it has little direct effect on consumption. Unsurprisingly, since our formula is designed to capture only causal effects, it does not match the simple correlation between consumption and house prices when house prices are driven by structural shocks, such as income, which more directly affect consumption.

2 Model

We consider a dynamic, incomplete markets model of household consumption. Households have finite lives and face uninsurable idiosyncratic income risk. The main distinguishing feature of the model is that households trade houses that provide housing services and can borrow against the value of their houses.

2.1 Set up

Time is discrete and runs forever. There is a constant population of overlapping generations of households, each living for J periods. The first J_y periods correspond to working age, the next J_o periods to retirement.

Households invest in two assets: a risk-free asset and housing. Let A_{it} and H_{it} denote the holdings of the two assets by household i at time t . The risk-free asset is perfectly liquid and yields a constant interest rate r . The housing stock yields housing services one-for-one, depreciates at rate δ , and trades at the price P_t . For most of our analysis we focus on the case of constant house prices, $P_t = P$, and study the effects of an unanticipated, permanent shock to P . However, in Appendix 3, we also consider the case in which P_t is stochastic and follows a random walk and, in Section 5, the case in which the expected price path is deterministic but not constant.¹⁸ In both cases results are extremely similar to those obtained with the simpler one-time unanticipated shock.

Households born at time t maximize the expected utility function

$$E \left[\sum_{j=1}^J \beta^j U(C_{it+j}, H_{it+j}) + \beta^{J+1} B(\tilde{W}_{it+J+1}) \right],$$

where C_{it} is consumption of non-durable goods and \tilde{W}_{it+J+1} is the offspring's real wealth. The per-period utility function and the bequest function are, respectively,

$$U(C_{it}, H_{it}) = \frac{1}{1-\sigma} (C_{it}^\alpha H_{it}^{1-\alpha})^{1-\sigma}, \quad B(\tilde{W}_{it+J+1}) = \Psi \frac{1}{1-\sigma} \tilde{W}_{it+J+1}^{1-\sigma}.$$

The offspring's real wealth is

$$\tilde{W}_{it+J+1} = \frac{P_{t+J+1}(1-\delta)H_{it+J} + (1+r)A_{it+J}}{P_{Xt+J+1}},$$

where $P_{t+J+1}(1-\delta)H_{it+J} + (1+r)A_{it+J}$ are bequests, and $P_{Xt+J+1} = \Omega P_{t+J+1}^{1-\alpha}$ is a price index

¹⁸The focus on permanent shocks has a good empirical motivation. Real house prices exhibit values of annual persistence around 0.94 in Case-Shiller and OFHEO house price data from 1960-2014, and a random walk cannot be rejected.

that adjusts for changes in the cost of housing.¹⁹

The assumption of Cobb-Douglas preferences for non-durable consumption and housing services plays an important role in our baseline results. While estimates of the elasticity of substitution between non-durables and housing based on macro data are somewhat varied, more relevant evidence from micro data consistently finds support for an elasticity of substitution close to unity (Piazzesi et al. (2007), Davis and Ortalo-Magné (2011), and Aguiar and Hurst (2013)). Furthermore, in Appendix 2 we obtain similar quantitative results using CES preferences with elasticity of substitution in a reasonable range.

Households face an exogenous income process. When the household works, income is given by

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

where $\chi(j_{it})$ is a deterministic age-dependent parameter, j_{it} is the age of household i at time t , and z_{it} is a transitory shock that follows an AR(1) 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, which we specify following Guvenen and Smith (2014).

The per-period budget constraint is

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

Households can borrow, but they have to satisfy the borrowing constraint

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

where $(1 - \theta)$ is the fraction of a house's future value that can be used as collateral.

2.2 The Permanent-Income Case

We begin with a special case that can be solved analytically and gives us a reference permanent-income-hypothesis (PIH) result, with small house price effects on consumption.

Consider the case of a deterministic income process ($\sigma_\varepsilon = 0$), no borrowing constraints, and constant house prices $P_t = P$. Assume:

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

¹⁹In particular, $\Omega = \alpha^{-\alpha}(1 - \alpha)^{-(1-\alpha)}(1 - \frac{(1-\theta)(1-\delta)e^g}{1+r})$

Under these assumptions, there is perfect consumption smoothing: consumption of non-durable goods and housing are constant over the lifetime. Moreover, non-durable consumption is equal to a fixed fraction $\alpha(1 - \beta)$ of total wealth, which includes human wealth, housing wealth, and financial wealth:²⁰

$$C_{it} = \alpha(1 - \beta) \left[\sum_{\tau=t}^{t+J-j} (1+r)^{-\tau} Y_{it+\tau} + (1-\delta) PH_{it-1} + (1+r) A_{it-1} \right].$$

Now consider the effect of an unexpected, permanent shock to the house price P . The elasticity of consumption to this shock is equal to the share of housing in total wealth:

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

What are the quantitative implications of this case? To assess this, note that equation (2) also holds in the aggregate and thus each quantity on the right-hand side can be measured directly using aggregate data. For consistency with the rest of the paper let us use aggregates from the 2001 Survey of Consumer Finances (SCF). We then get $(1 - \delta)PH = 2.15Y$ and $A = -.32Y$, where H is average housing, A is average liquid wealth net of debt, and Y is average earnings. Using an interest rate of $r = 2.5\%$ and an infinite horizon approximation, human wealth is equal to $Y/r = 40Y$. The aggregate elasticity of non-durable consumption implied by the model is then 0.0514, a small number relative to empirical housing price effects. Moreover, this elasticity is insensitive to the level of indebtedness of the household sector. For example, suppose that household debt goes up by $0.5Y$ so that $A = -0.82Y$, holding all else equal. This is a very large increase in debt but yields a nearly identical and still small elasticity of 0.0520. Finally, this simple model implies that elasticities will be higher for older households, since they have a smaller fraction of human wealth to total wealth.

What drives the consumption response to house prices in the PIH model? The response can be decomposed into three effects: a substitution effect, an income effect, and an endowment effect.²¹ It is then possible to interpret equation (2) in two ways.

First, due to the Cobb-Douglas assumption, the income and substitution effects exactly cancel out. Since only the endowment effect remains, this implies that the change in consumption in (2) can be interpreted as a pure endowment effect.

However, an alternative interpretation is possible. In this model consumption of housing

²⁰For simplicity, we set the human wealth of offspring to zero.

²¹A permanent increase in P increases 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, keeping the present value of future expenditures constant. The income effect is the change in current consumption due to a reduction in the present value of expenditures arising from the increased 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 expenditures arising from the increase in the value of the initial housing stock.

services is constant over time. Hence, at any point after the first period of life, an increase in the price of housing raises the value of an agent's housing endowment, but at the same time it raises the net present value of the future implicit rental cost on housing services by roughly the same amount.²² The detailed derivations behind these statements are in Appendix 1. Therefore, the effect in (2) can also be interpreted as an (almost) pure substitution effect, with the income and endowment effects canceling out. This interpretation is consistent with the view discussed in the introduction that housing price effects must be small, because of the increase in future implicit rental costs. It is important to note that both interpretations of (2) are correct. However, the first interpretation will be especially useful in what follows as it survives in richer versions of the model.

2.3 Sufficient-Statistic Formula

Return now to the general model with income uncertainty and borrowing constraints, and assume house prices follow the geometric random walk with drift

$$P_t = x_t P_{t-1},$$

where x_t is an i.i.d. shock with $E[x_t] = e^g$. In this setup, we derive our main analytical result: the individual consumption response to a permanent house price shock is given by a simple formula, the individual marginal propensity to consume out of temporary income shocks times the value of the housing stock.

To set the stage for the result, we first represent the household problem recursively. Define total wealth $W_{it} \equiv (1 - \delta)P_t H_{it-1} + (1 + r)A_{it-1}$. The household's Bellman equation can be written as

$$V(W_{it}, z_{it}, j_{it}, P_t) = \max_{C_{it}, H_{it}, A_{it}} U(C_{it}, H_{it}) + \beta E[V(W_{it+1}, z_{it+1}, j_{it} + 1, x_{t+1}P_t)] \quad (3)$$

$$\begin{aligned} \text{subject to } \quad & C_{it} + P_t H_{it} + A_{it} = Y_{it} + W_{it}, \\ & W_{it+1} = (1 - \delta)x_{t+1}P_t H_{it} + (1 + r)A_{it} \quad \forall x_{t+1}, \\ & (1 - \theta)(1 - \delta)x_{t+1}P_t H_{it} + (1 + r)A_{it} \geq 0 \quad \forall x_{t+1}. \end{aligned}$$

The bequest motive gives the terminal condition:

$$V(W_{it}, z_{it}, J + 1, P_t) = \frac{\Psi}{1 - \sigma} \left(\frac{W_{it}}{P_{Xt}} \right)^{1 - \sigma}.$$

²²The effects are not exactly equal due to depreciation δ . When $\delta = 0$ they are exactly equal.

We are now ready to prove our main analytical result.

Proposition 1 *The individual response of non-durable consumption to the permanent house price shock x_t is*

$$MPC_{it} \times (1 - \delta) P_{t-1} H_{it-1}, \quad (4)$$

where MPC_{it} is the individual marginal propensity to consume out of transitory income shocks.

Proof. First, we prove by induction that the value function can be written as

$$V(W, z, j, P) = P^{-(1-\sigma)(1-\alpha)} v(W, z, j),$$

for all W, z, j, P , where for $j < J + 1$ the function v satisfies the Bellman equation

$$v(W, z, j) = \max_{C, \tilde{H}, A, W'} U(C, \tilde{H}) + \beta E [x^{-(1-\sigma)(1-\alpha)} v(W', z', j + 1)],$$

$$\begin{aligned} \text{subject to} \quad C + \tilde{H} + A &= Y(s) + W, \\ W' &= (1 - \delta) x \tilde{H} + (1 + r) A, \\ (1 - \theta) (1 - \delta) x \tilde{H} + (1 + r) A &\geq 0, \\ z' &= \rho z + \epsilon. \end{aligned}$$

The property holds immediately for $V(W, z, J + 1, P)$ since $P_{Xt} = \Omega P_t^{1-\alpha}$. Next, we prove that if the property holds for $V(W, z, j + 1, P)$ then it holds for $V(W, z, j, P)$. Rewrite the Bellman equation (3) in terms of the variable $\tilde{H}_t = P_t H_t$. The property $U(C, \tilde{H}/P) = P^{-(1-\sigma)(1-\alpha)} U(C, \tilde{H})$ and the induction hypothesis imply that the expression $P_t^{-(1-\sigma)(1-\alpha)}$ can be factored out of the objective function and does not affect the optimization problem. This completes the induction step.

Let $C(W, z, j)$ denote the policy function associated with this optimization problem and notice that it is independent of the current price P . Therefore, C_{it} responds to the shock x_t only through its effect on W_{it} and the response is

$$\frac{\partial C(W_{it}, z_{it}, j_{it})}{\partial W} (1 - \delta) P_{t-1} H_{it-1}.$$

To complete the argument notice that $\partial C(W_{it}, z_{it}, j_{it}) / \partial W$ is also equal to the response to a transitory income shock, denoted by MPC_{it} . ■

The result can immediately be restated in terms of the individual elasticity of non-durable consumption to house prices

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

The consumption response to house price changes can be measured in different ways. The elasticity η_{it} has the advantage of being unit free. The expression (4), which captures the response in dollar terms to a percentage increase in house values has the advantage that it can be aggregated over individuals to yield the economy-wide response of consumption to a uniform percentage change in house values. Finally the consumption response can be measured in dollar-per-dollar terms, as the response of consumption to a one dollar increase in house values, in which case Proposition 1 states that the response is equal to MPC_{it} . In the rest of the paper, we refer to all these three measures, but tend to favor the elasticity.

Notice that Proposition 1 can be easily extended to characterize the response to any permanent component of price shocks in models with general stochastic processes for P_t . The responses to transitory shocks and to shocks that affect the expected future growth rate of house prices will be discussed in Section 5.

The simple formula in Proposition 1 contains two objects that are both endogenous, so it does not deliver a closed form solution. However, the formula is useful for understanding how endogenous forces determine the sensitivity of an economy to house price shocks. In particular, house price shocks will have bigger effects when MPCs are larger, when gross housing wealth is larger, and when there is a stronger correlation between MPCs and housing values in the economy.

In Section 3, we show that the formula continues to hold approximately in richer versions of the model with adjustment costs, more realistic mortgages and the option to rent. Hence, it can provide intuition in substantially more complicated environments. Furthermore, given empirical estimates of MPCs conditional on housing wealth, the formula can be used as a sufficient statistic for the model-implied response of consumption to house price shocks, without the need to structurally estimate the full model. We explore this approach in Section 4.

At first sight, formula (4) may appear tautological. The non-obvious content arises because we are considering the consumption response to a change in house prices not endowments of housing wealth. A house price shock changes current endowments, but *also* changes relative prices and the tightness of borrowing constraints so it is not trivial that responses would depend only the endowment effect which arises in the current budget. To better understand the result it is useful to think about the underlying forces at work. As previewed in the introduction, when house prices increase, there are four effects: (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 poorer overall because the implicit rental cost of housing is higher today and in all future dates; (3) a collateral effect, arising from the fact that higher collateral values allow households to increase borrowing, today and in the future; (4) an endowment effect, which comes directly from the increase in the value of the house owned when the shock hits.²³ Formula (4) tells us that the first three effects exactly cancel out, so only the endowment effect

²³To formally define income and substitution effects requires extending these notions to the dynamic, incomplete markets environment analyzed here. The details are presented in Appendix 1.

remains. We will return to this decomposition in Subsection 2.6, after calibrating and simulating our model, to look at the relative magnitude of these four effects.

The Cobb-Douglas assumption on preferences clearly plays an important role in the proof of Proposition 1. In Appendix 2, we examine the case of general CES preferences and show that our main quantitative conclusions survive as long as the elasticity of substitution between consumption and housing is not too far from one.

2.4 Calibration

We now calibrate the model in order to assess its quantitative implications for consumption responses to house prices, and we use Proposition 1 to interpret the results.

The baseline model parameters are shown in Table 1. The model is annual. We interpret the first period of life as age 25.²⁴ 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 $r = 2.4\%$.²⁵ In our baseline calibration, we use a coefficient of relative risk aversion σ equal to 2.²⁶

House prices are constant except for an unexpected, permanent shock. We choose a depreciation rate of housing $\delta = 2.2\%$ to match the depreciation rate in BEA data from 1960-2014. The collateral constraint parameter θ determines the minimum mortgage down payment, and we choose a conservative value of 0.25.²⁷

Table 1: *Parameter Values*

Parameter	σ	r	δ	θ	ρ_z	σ_z	α	β	Ψ	Ξ
Value	2	2.4%	2.2%	0.25	0.91	0.21	0.854	0.9331	1,353	3.4137

The income process during working age has a life-cycle component and a transitory component. The life-cycle component is chosen to fit a quadratic regression of yearly earnings on age in the PSID as in Kaplan and Violante (2010). Following Floden and Lindé (2001), the temporary component z follows an AR1 process with autocorrelation $\rho_z = 0.91$ and standard deviation $\sigma_z = 0.21$ to match the same statistics in PSID earnings (after taking out life-cycle components). In retirement, households receive a social security income payment which is modeled as in Guvenen and Smith (2014).

The three remaining parameters to calibrate are: the coefficient α on housing services, the discount factor β , and the bequest parameter Ψ . We choose these parameters jointly to match

²⁴We begin the model at this age in order to abstract from complications with schooling decisions.

²⁵We target interest rates from 1990-2000, which we view as the steady-state for our simulations but results are not sensitive to the level r .

²⁶Using log utility or other values of σ does not substantively change our conclusions on the effect of a price level shock. In Section 5 we discuss how σ affects the model response to anticipated price changes.

²⁷Lowering θ amplifies the size of consumption responses.

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).²⁸ Our model is very stylized for retired agents since they face no sources of risk. Therefore, we prefer to focus our calibration and predictions on working-age agents. Our notion of liquid wealth net of debt in the SCF excludes retirement accounts for agents before retirement, but includes retirement accounts for agents above age 60. In the model, we assume that retirement accounts take the form of a lump sum transfer at retirement, which is calibrated to equal a fraction Ξ of labor income prior to retirement.

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.²⁹ We then 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 target liquid wealth rather than total non-housing wealth because this delivers MPCs which are more in line with empirical estimates. 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 withdrawal penalties prior to retirement but become fully liquid after retirement. Thus, we believe that our calibration strategy reasonably matches the fraction of wealth that can be easily accessed both prior to and after retirement.³⁰

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.³¹ The blue circles are the model predictions and the green squares are from the 2001 SCF. Despite its simplicity the model delivers a reasonable fit, the main discrepancy being too little housing in the early periods and too much debt in the mid 30s and 40s.

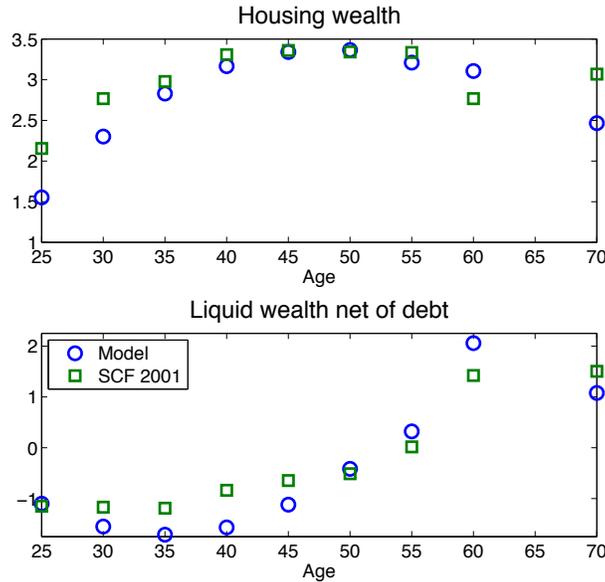
²⁸More specifically, following Kaplan and Violante (2014), we define liquid wealth net of debt in the data as the sum of cash, money market, checking, savings, and call accounts as well as directly held mutual funds, stocks, bonds and T-bills, net of credit card and mortgage debt and pick this to match net liquid assets in the model. For retired households we also include retirement accounts in liquid assets. Note that we subtract mortgage debt since households can borrow frictionlessly against housing in the model, so a household with the same financial assets but less mortgage debt has more ability to smooth consumption. Housing wealth in the data is defined as the gross value of primary residences, other residential real estate and nonresidential real estate, and is matched to H in the model. In order to normalize the data and model, we divide liquid assets net of debt and housing wealth by the overall average income of age 25-59 year old households.

²⁹Using only age 25 households would give us a small sample and introduce large measurement error.

³⁰Nevertheless, targeting total wealth rather than liquid wealth still produces large elasticities.

³¹The point for the age bin 65 and over is plotted at age 70.

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



2.5 The effects of a permanent house price shock

We now turn to the model predictions for our main comparative static exercise. Namely, we look at the instantaneous response of non-durable consumption to a negative, unexpected, permanent 1% reduction in house prices.³²

The average elasticity of consumption over working life is 0.47. Figure 2 reports elasticities for different ages, 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 main takeaway is that magnitudes are much larger than in the PIH model and are declining rather than increasing with age. In fact, the elasticities obtained here are considerably larger than empirical estimates in the literature. In the following section, we show that adding more realistic features to the model, and in particular adding the option to rent, yields elasticities more in line with empirical estimates.

We can use our formula (4) to interpret this result. In Figure 3 we separately plot the two elements of the formula: MPCs and housing values, averaged by age. The high elasticities of young agents are driven by the fact that they have a high MPC and, at the same time, hold substantial amounts of housing. In a precautionary saving model the MPC depends on total net worth W and is decreasing in W due to concavity of the consumption function. Young agents own houses but finance them with debt, so they have low net worth despite having relatively high levels of housing. This explains the joint presence of high MPCs and high levels of housing.

³²Results are very similar for larger house price declines as well as house price increases. Our model does not deliver the asymmetries to house price increases and decreases implied by Guerrieri and Iacoviello (2014)'s representative agent model because whether or not borrowing constraints bind in our model is largely determined by idiosyncratic shocks rather than aggregate house price movements.

Figure 2: Elasticities over the Life Cycle: Frictionless Model

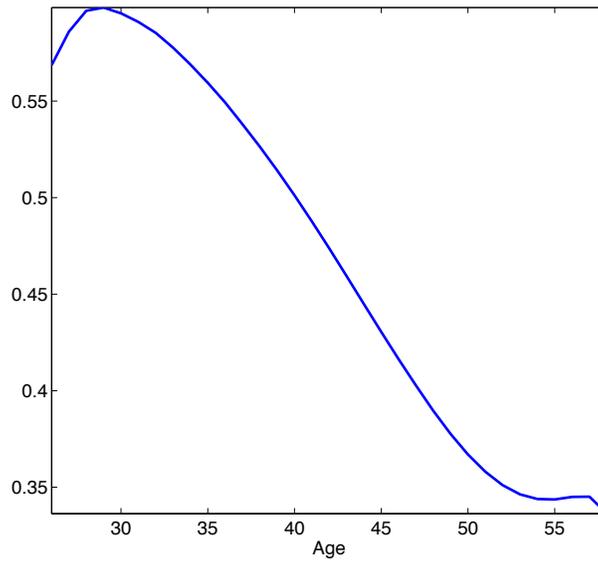
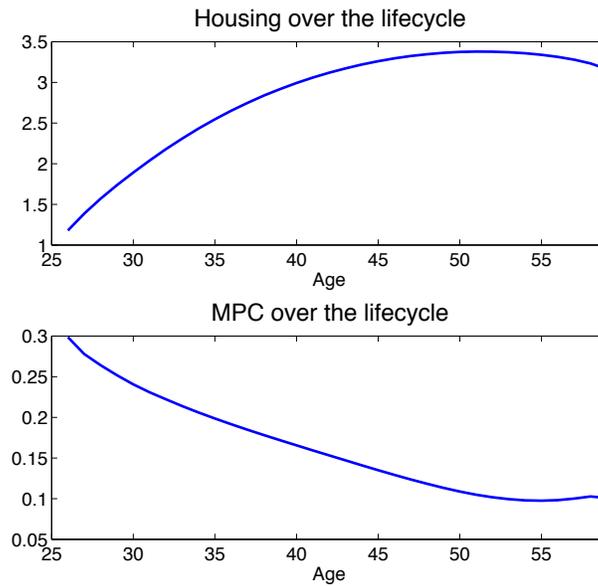
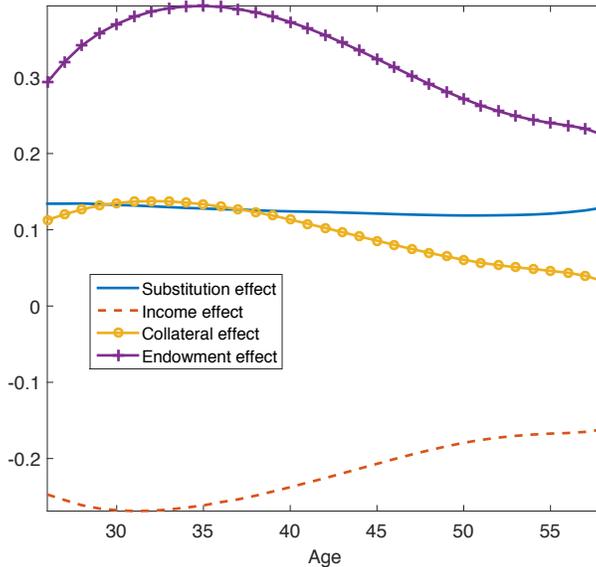


Figure 3: Understanding the Life Cycle: Frictionless Model



Interestingly, our finding that young homeowners, who are more levered, exhibit larger responses to house prices than old homeowners, who are less levered, is consistent with the empirical finding in Attanasio et al. (2009).

Figure 4: Decomposition of The House Price Effect



2.6 Decomposing The House Price Effect

To conclude this section, we return to the decomposition of the consumption response introduced in Subsection 2.3 and discuss its connection with the existing literature. As mentioned before, the response can be decomposed into four effects: substitution, income, collateral, and endowment. Proposition 1 shows that the first three effects exactly cancel and the response of consumption to a house price shock can be interpreted as purely driven by the endowment effect. However, this does not imply that each of the other effects is individually small in absolute value. In Appendix 1 we develop a methodology to distinguish the four effects, extending definitions from classical demand theory to models with uninsurable risk. The effects are reported by age in Figure 4. For this particular figure we report the response of the consumption level to a 1% change in house prices rather than elasticities because the decomposition is additive in levels and so aggregates easily across agents.

One can see from Figure 4 that the income effect cancels with the sum of the substitution and collateral effects. This implies that, as shown by our sufficient-statistic formula, the total consumption response can be interpreted as a pure endowment effect. However, an alternative interpretation of the same figure is that the total consumption response reflects the sum of three effects which are each important and similar in magnitude: the net wealth effect, which is given by the endowment effect minus the income effect, the substitution effect, and the collateral effect.

Both interpretations are correct and useful for different reasons: The latter interpretation shows that many competing effects are at work, and each is relevant on its own for determining the total consumption response. The former interpretation shows that these effects are tightly

linked together through optimizing behavior and that the net effect can be summarized by the product of two easily interpretable empirical objects: the MPC and the house value.

To illustrate the usefulness of both interpretations, consider the large consumption response in our model relative to the PIH model. Part of the large response in our model directly reflects the addition of the collateral channel. In addition, unlike in the PIH model, the endowment effect no longer approximately cancels with the income effect. Instead, the net wealth effect of a house price increase is positive and sizable. The intuition for this result is that a permanent increase in house prices leads to an immediate positive endowment effect while the increase in implicit rental costs occurs mostly in the future. Even though house price movements remain roughly neutral for a household’s lifetime budget constraint, borrowing constraints mean that consumption responds more to current than future income, so the endowment effect dominates. Thus, the greater consumption response in our model relative to the PIH model can be interpreted as an increase in the net wealth effect, with an additional collateral channel. Alternatively, using our sufficient-statistic interpretation, it is immediate to see that the small consumption response in the PIH model simply reflects the fact that this model implies a very small MPC.

The fact that borrowing constraints change the strength of the net wealth effect is also important when using PIH intuition to empirically identify house price shocks on consumption. For example, Attanasio et al. (2009) find that consumption of the young is more correlated with house prices than that of the old and interpret this as evidence against housing wealth shocks, since in a PIH framework, older households should respond more strongly to changes in wealth due to their shorter planning horizon. Attanasio et al. (2011) validate this prediction in a sophisticated quantitative model which relaxes the stylized PIH assumptions, but our results show that this theoretical relationship does not hold in general and can be somewhat sensitive to particular modeling choices. In our baseline results, the young respond much more strongly to house price shocks. In the following section, we allow for rental markets and find that consumption responses become largely independent of age. Our theoretical results show that the age-profile of consumption responses will be determined by the age-profile of $MPC \times PH$, which our quantitative results show is theoretically ambiguous.³³ Ultimately, this age-profile is an empirical question, which we address in Section 4.

Our theoretical decomposition of housing price effects also helps reconcile our large elasticities with the intuition of the “small house-price-effects” view mentioned in the introduction. This view builds on the PIH model by arguing that, since changes in housing asset values are offset by future changes in the user-cost of housing, “housing wealth is not real wealth” and should have little effect on consumption. For example, Sinai and Souleles (2005) argue that this should be true when house price movements are perfectly correlated since everyone must live somewhere. However, in their model there is no precautionary motive, which implies that income and endowment effects

³³This theoretical ambiguity might also help explain why some other studies such as Kaplan et al. (2015) and Campbell and Cocco (2007) find stronger effects for the old.

exactly cancel; there are no substitution effects, as housing holdings are fixed; and there are no collateral effects. Therefore, the three channels identified above are absent.

Finally, our results have implications for the empirical literature measuring the strength of the various channels through which house prices affect consumption. Each channel in our model is large in isolation, so our model is quite consistent with the conclusions in DeFusco (2015) and Mian et al. (2013) that there is a significant consumption response to housing collateral changes and that the poorest and most indebted households respond most significantly to house price movements through this channel. In particular, the collateral channel in our model implies a marginal propensity to borrow of just over 10 cents, which is in line with the evidence in DeFusco (2015). Nevertheless, it is important to note that the collateral effect in our model is tightly linked to the other three effects, and that their sum eventually boils down to the endowment effect alone, so the interpretation of the collateral effect in isolation must be done with care.

3 Extended Model

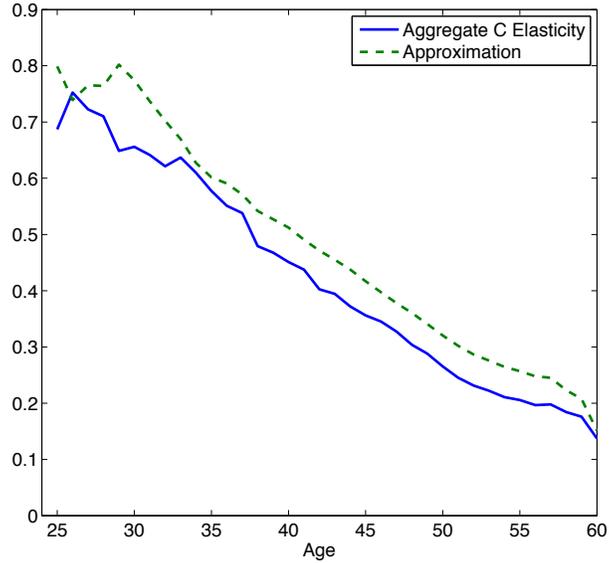
In this section, we extend the model by relaxing various simplifying assumptions in order to get a better quantitative assessment of housing price effects. In particular, we progressively enrich the model by introducing adjustment costs to housing, the option to rent, and more realistic mortgage contracts. While models with these features are now prevalent in the literature, the precise ways in which they shape housing price effects has largely been unstudied.³⁴ Thus, as with the rest of the paper, the goal is not to develop a novel model of consumption but is instead to characterize the previously unrecognized common features in existing models which determine consumption responses to house price movements.

In terms of our sufficient-statistic approach, our main conclusion is that the formula derived in Proposition 1 provides a good approximation for all realistic calibrations explored in this section. This means that the insights developed in the previous section carry forward and can continue to provide intuition for the workings of these progressively more complicated models. In essence, all of the complicated features of these models can still be summarized by their effects on our simple sufficient-statistic. This does not mean that the quantitative implications are insensitive to the different specifications analyzed, since different specifications have different implications for the joint distribution of housing wealth and MPCs, but it does mean we can use the formula to interpret results obtained under all of our calibrations.

In terms of magnitudes, we still find fairly large elasticities under all the calibrations explored. The fully enriched model produces elasticities that are roughly half of what we found in the baseline, and is thus closer to the empirical estimates. This is mostly due to the introduction of the rental option, as we explain shortly.

³⁴See Cocco (2005) and Yao and Zhang (2005) for two of the earliest papers incorporating these elements.

Figure 5: Elasticities over the Life Cycle: Model with Adjustment Costs



3.1 Housing Transaction Costs

The assumption that housing can be traded frictionlessly is clearly counterfactual: search costs and brokers' fees make it costly to sell and buy houses, and indeed households trade houses only infrequently.³⁵ We model adjustment costs as a fixed cost, incurred whenever the household changes its stock of housing. In particular, if household i decides to trade housing at time t , it pays a cost proportional to the value of the house sold

$$\kappa_{it} = F \cdot P_t H_{it-1} \mathbb{1}_{H_{it} \neq H_{it-1}},$$

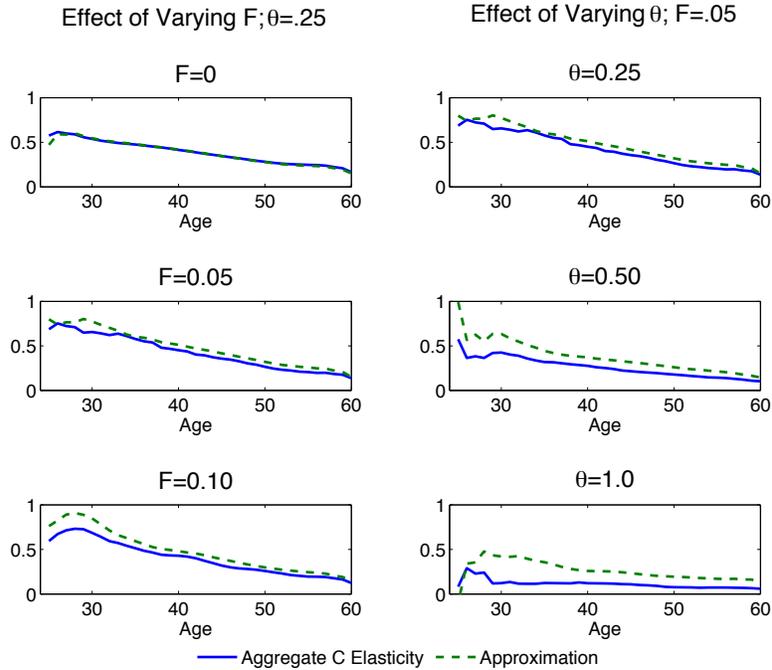
where $\mathbb{1}$ is an indicator function equal to 1 iff $H_{it} \neq H_{it-1}$ and where we set $F = 0.05$. This transaction cost is equal to the value of housing adjustment costs calibrated in Díaz and Luengo-Prado (2010) and is close to the adjustment costs of 0.0525 estimated in Berger and Vavra (2015) for a broad measure of durable spending. In addition, we show below that our conclusions are not particularly sensitive to changes in the size of this cost.

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. In Appendix 4, we provide a detailed description of the numerical solution method. We follow the same calibration strategy as in the baseline and report the parameters in Appendix 5.

Once we introduce transaction costs, Proposition 1 no longer holds. In Figure 5 the solid

³⁵Berger and Vavra (2015) show that the average annual frequency of housing adjustment in the PSID is 4.3% in data from 1968-1996 and 5.8% in data from 1999-2011.

Figure 6: Robustness to Different Fixed Costs and Downpayment Parameters



blue line shows the computed average elasticity of non-durable consumption as a function of age and the dashed green line shows the sufficient-statistic-implied elasticity from Proposition 1. The figure shows that even though the formula does not give the exact elasticity anymore, it remains a quite accurate approximation.

The shape of the average elasticity over the life cycle is similar to the one in the baseline model: it is high for the young agents who have high MPCs as well as substantial housing and falls as households age and accumulate assets. Households can borrow to 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 when young, using leverage, rather than slowly accumulating housing. This explains why the elasticities for young households are larger in Figure 5 than in Figure 2.

How robust is the accuracy of the approximation formula in Proposition 1 to alternative environments? Figure 6 compares the computed average elasticity with the sufficient-statistic-implied elasticity for different values of the fixed cost F and of the downpayment parameter θ . The panels on the left keep θ at its baseline value and change F , the panels on the right keep F at its baseline value and change θ .

The figure shows that the formula of Proposition 1 is an accurate approximation for a range of values of F and θ . In particular, increasing F barely affects the accuracy of the approximation as long as θ remains low. The intuition is that, as long as households can borrow against their house values, consumption can respond even if transaction costs substantially reduce house trading. The consumption response becomes much more muted than what is implied by the $MPC \times PH$

formula only when unrealistically high values of F and θ are combined, so that households can no longer respond to the increased house price by either selling or borrowing.

3.2 Rental Markets

Roughly one third of US households rent rather than owning, and their (non-existent) housing wealth is unaffected by changes in house prices. Therefore, it is important to introduce a rental option. To do so, we assume that in each period a household must choose whether to be a homeowner or a renter. Renters pay a flow rental cost R_t per unit of housing services. Rental housing can be adjusted costlessly but cannot be used as collateral. We assume that the rental cost R_t is proportional to house prices, $R_t = \phi P_t$, so that house price movements are passed proportionally into rental costs. That is, in our baseline results, we assume a constant price-rent ratio. In Appendix 5, we show that elasticities are mildly amplified if we instead make the opposite extreme assumption that R_t is fixed while P_t varies, so that our baseline choice is relatively conservative.

Households face a trade-off when choosing between renting and owning. The advantage of renting is that it allows the household to keep its savings in the form of liquid assets, thus providing a better buffer against income shocks. The disadvantage is that renting is costlier, as we assume $\phi > (r + \delta)/(1 + r)$,³⁶ and rental housing cannot be used as collateral.³⁷ Young households who are more financially constrained are the ones who benefit more from the flexibility of renting and so will disproportionately choose to rent.

The rental-to-price ratio ϕ is an additional parameter of the model. We follow the previous calibration strategy for all other parameters but also target the life-cycle profile of homeownership rates.³⁸ All parameters are reported in Appendix 5. Figure 7 shows that we are still able to match the life-cycle profile of housing and liquid wealth reasonably well. In addition, the model can fairly well match the upward slope and subsequent flattening of homeownership rates over the life cycle.

How does the addition of the rental option affect the size of consumption responses to house prices and the accuracy of our approximation? Figure 8 shows the actual average elasticity over the life cycle for the model with and without the rental option (solid blue line and solid red line). The figure also shows the sufficient-statistic approximation for both models (dashed blue line

³⁶Notice that $(r + \delta)/(1 + r)P$ is the user cost of housing services for a homeowner facing a constant house price P .

³⁷The latter advantage is not enough, on its own, to induce households to own, because renting reduces the need for borrowing more than it reduces the availability of collateral, since $\theta < 1$.

³⁸At retirement, income risk falls to zero. This substantially changes the trade-off between liquid and illiquid assets. With a constant rental rate this would imply a large jump up in the homeownership rate at retirement. To eliminate this jump, we introduce a different rental-to-price ratio at retirement $\hat{\phi} < \phi$, and we pick this additional parameter to match homeownership rates after retirement. 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. At any rate, since we concentrate on results for working-age households, the choice of $\hat{\phi}$ has little effect on any of our results.

Figure 7: Life-Cycle Moments: Data vs Model

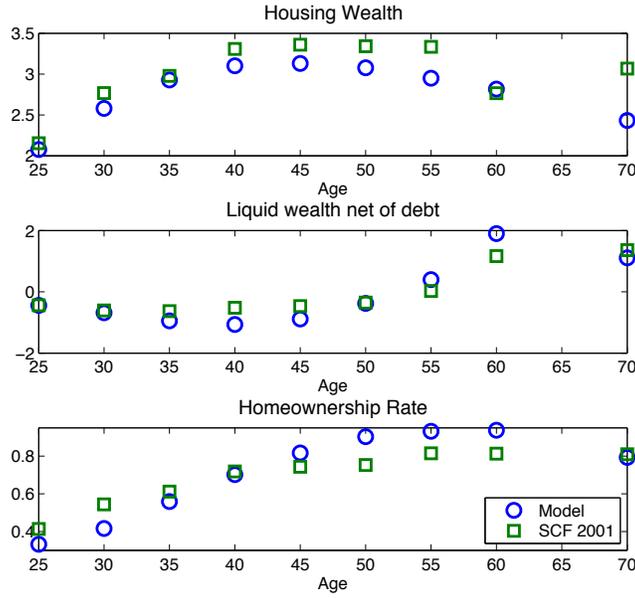
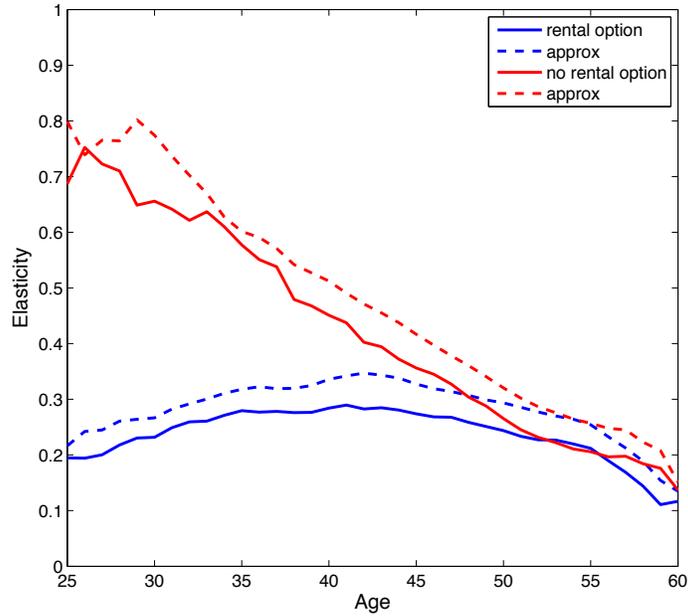


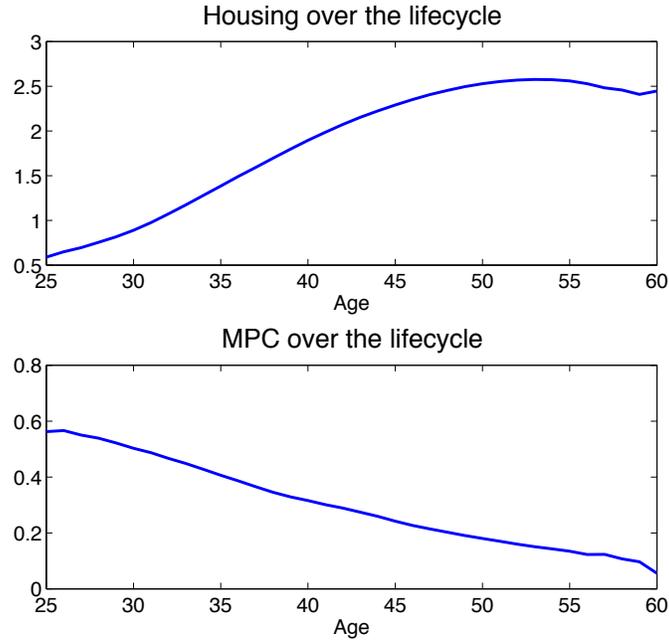
Figure 8: Elasticities over the Life Cycle



and dashed red line). There are two takeaways from this figure. The first is that even with the rental option, the sufficient-statistic formula of Proposition 1 remains a good approximation of the true model elasticity.³⁹ The second is that the addition of a rental option substantially lowers

³⁹This sufficient-statistic is also accurate separating homeowners and renters. Renters' true elasticity is zero, as is also implied by the sufficient-statistic formula since $MPC = 0$. While one might expect that decreases in house

Figure 9: Understanding the Life Cycle



the response of consumption to house prices, especially for the young households. The average elasticity over working life goes from 0.43 in the model without the rental option to 0.24 in the model with rentals.⁴⁰ Therefore, the introduction of rentals brings us closer to the estimates from the empirical literature discussed in the introduction.

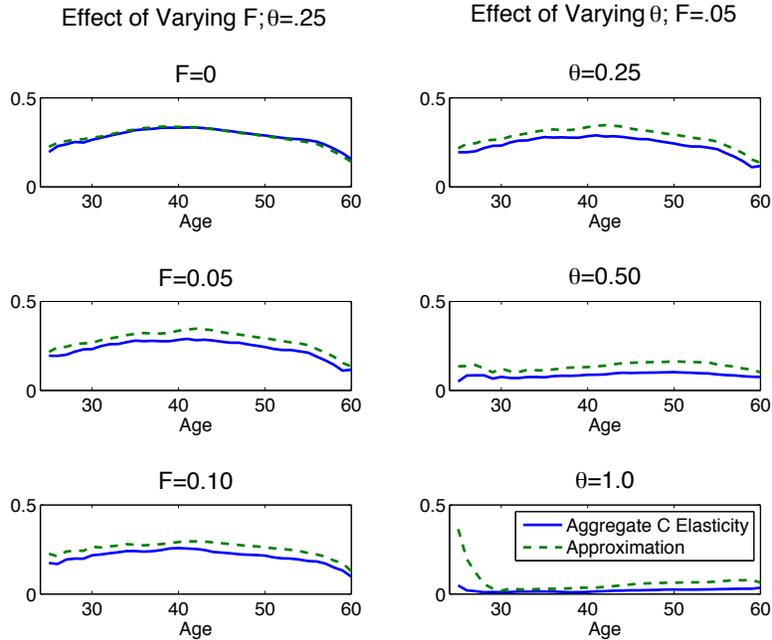
Our $MPC \times PH$ formula helps to provide intuition for why the elasticity falls when households have the option to rent. First, since renters do not own a house, their consumption does not respond to house price movements. Hence, if we match the homeownership rate in our dataset of roughly 70%, the average elasticity of consumption to house prices will mechanically fall, even if all households are otherwise identical. However, not all households are identical and home ownership is not randomly distributed in the population. Figure 9 shows the life-cycle profile of housing and MPCs in the model with the option to rent. Notice that the young households with high MPCs are the ones who choose to rent and thus have low levels of housing.⁴¹ This in turn lowers their consumption elasticity. Furthermore, within every age group, households with high MPCs tend to choose to rent, given that they are closer to being constrained and thus prefer liquid wealth over illiquid housing. Thus, our model rationalizes the empirical findings in

prices would lead to increases in renters' consumption since rental prices also decline, this is not the case in this class of models. Renters own no housing and so their endowment and collateral effects are zero. The decrease in the price of housing induces a positive income effect on consumption but this is offset by the substitution effect as households switch from housing to consumption.

⁴⁰If we include retired households average elasticities fall from 0.40 to 0.21 when introducing rental markets.

⁴¹Indeed, in both the model and the data, the life-cycle profile of homeownership is steeply upward sloping. By age 45 the models with and without rental markets are nearly identical.

Figure 10: Robustness to Different Downpayments and Fixed Costs



Campbell and Cocco (2007) and Guiso et al. (2005) that renters of all ages exhibit small responses of consumption to house price movements.

Finally, Figure 10 assesses the accuracy of the sufficient-statistic formula under a range of alternative parameter values. As in the model with no rental option, the accuracy of the approximation 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 still quite accurate for reasonable values of this parameter. In particular, as θ rises, the elasticity falls. This is now the combination of two factors. First, household leverage falls, which reduces MPCs for a given housing value. More importantly, in these comparative statics exercises we do not recalibrate the model. Thus, as we raise the required down payment while keeping ϕ constant, home ownership 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, the homeownership rate in the model falls to 23%. Therefore, most of the decline in the elasticity as θ increases occurs because the homeownership rate is driven to counterfactually low levels.

3.3 More realistic mortgages

So far we have modeled borrowing as one-period loans subject to a collateral constraint. We now briefly discuss extensions that more realistically capture how households borrow against their houses' value. While we do not attempt to model all the details of fixed-rate mortgages and home-equity loans, we can easily increase the model's realism by introducing costly equity

extraction together with long-term mortgages in which the collateral constraint is only binding at loan origination.⁴²

We have so far assumed that households can costlessly increase borrowing as long as they meet the collateral constraint. In reality, there are costs associated with extracting housing equity by refinancing. This means that even in the model with transaction costs on housing adjustment, our model probably overstates the liquidity of housing wealth. We now assess the robustness of our results to introducing a cost of “cash-out refinancing”. In particular, we extend the model by making the following assumption: if a household chooses positive holdings of the risk-free asset A_t or if it chooses any $A_t \geq A_{t-1}$ it can do so at no cost; but if the household chooses a negative A_t smaller than A_{t-1} —i.e., if it increases its debt level—it must pay a fixed cost.

Since the majority of refinancing costs do not depend on loan size, we model this cost as a fixed numeraire amount, in contrast to the housing transaction cost which is proportional to the size of the house.⁴³ We pick this cost of refinancing to match the fraction of refinancing observed in Bhutta and Keys (2014), which implies a fixed cost of 0.005.⁴⁴ In the model with no refinancing cost, roughly 30% of households increase debt each year while with the refinancing cost, this number falls to a realistic 10%.⁴⁵ Thus, refinancing costs substantially reduce the number of households who extract equity. However, the top panel of Figure 11, compared to Figure 8, shows that this modification has negligible effects on the size of elasticities.⁴⁶ The top panel of Figure 11 also shows our sufficient-statistic approximation remains accurate in this extension.

Mortgages are long-term loans. This introduces an asymmetry in the response of the borrowing limit to house price shocks. When house prices rise and households gain equity, households have the option to extract this equity and borrow more. In contrast, when house prices fall, lenders cannot force households to put up additional collateral. That is, households are not forced to repay their mortgage faster, even if their loan-to-value ratios go above the maximum allowed for new loans.⁴⁷ Since interest rates are constant in our model, long-term mortgages can be introduced with a simple change to the collateral constraint for households who are not adjusting their stock of housing. In particular, instead of (1) when not adjusting, we have the collateral

⁴²Note that since interest rates are fixed in our model, long-term mortgages in our environment are equivalent to the fixed rate mortgages prevalent in U.S. borrowing.

⁴³See Appendix 4 for more details. We have also solved versions of the model with refinancing costs proportional to the size of the house and results were extremely similar.

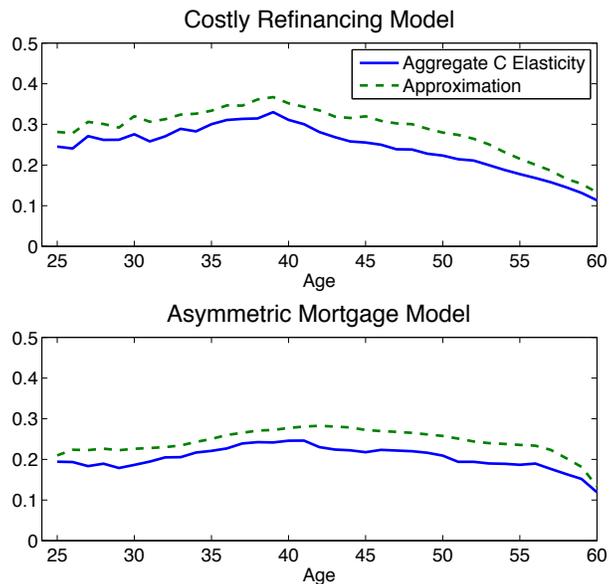
⁴⁴Since income is normalized to 1, this refinancing cost is 0.5% of average annual income, but even using much larger values produces similar results. For example, we have solved a version of the model where equity can only be extracted by buying/selling one’s house. Even in this extreme case, results are quite similar to Figure 11.

⁴⁵Our model has no distinction between forms of equity extraction, so we match total shares.

⁴⁶This is due to extensive margin effects which arise in environments with fixed adjustment costs. Adjustment costs reduce cash-out refinancing substantially, but this reduction is concentrated amongst households who would have only extracted small amounts of equity with small resulting effects on consumption. Any household who needs to extract substantial housing equity in order to increase consumption will still choose to do so even in the presence of adjustment costs.

⁴⁷In our model, given constant house prices, the maximum loan-to-value ratio is $(1 - \theta)(1 - \delta)/(1 + r)$.

Figure 11: Elasticities over the Life Cycle: More Realistic Mortgage Models



constraint:

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

The bottom panel of Figure 11 shows the elasticity to a 1% decline in house prices under this specification.⁴⁸ This variant leads to modestly lower, but quite similar elasticities to those in Figure 8. In addition, the sufficient-statistic approximation remains quite accurate. If house prices instead increase, then elasticities are identical to those in the baseline model since the modified collateral constraint only affects decisions when households are underwater. Thus, the introduction of long-term mortgages induces a small asymmetry in which price increases have slightly larger effects than price decreases, but this effect is quite modest.

4 Sufficient Statistic: Empirical Evidence

The previous sections show that a simple sufficient-statistic formula provides a good approximation for the consumption response to a permanent house price shock for a wide class of life-cycle models with uninsurable idiosyncratic income risk and borrowing constraints. In this section, we attempt to directly estimate the sufficient-statistic formula using micro data. This provides a new measure of housing price effects that does not rely on the specific parameters used to calibrate the model, but only on the model's general structure. We then show that our estimate is not far from existing measures obtained with different approaches.

⁴⁸We have also experimented with a larger price drop of 10% and the effect of the modified constraint is still small.

4.1 BPP approach

Estimating our sufficient statistic requires data on both the MPC out of a transitory income shock and on home values. While home values are easily obtained, estimating MPCs is more difficult. To estimate MPCs we follow the identification approach of Blundell et al. (2008) (henceforth BPP).

BPP show that if income follows a process with a permanent and an i.i.d. component, then, given individual level panel data on income and consumption, one can identify the MPC out of transitory shocks. In particular, assume that log income is $y_{it} = z_{it} + \varepsilon_{it}$, 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 $\Delta y_{it} = \eta_{it} + \Delta \varepsilon_{it}$.⁴⁹ Given this income process, the true MPC (in logs) out of a transitory shock is equal to⁵⁰

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

Under the assumption that households have no advanced information about future shocks, a consistent estimator of this MPC is

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

With panel data containing at least three time periods, one can implement this estimator with 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} .⁵¹ Since this requires individual level panel data on income and consumption, we use data from the Panel Study of Income Dynamics (PSID).⁵²

Kaplan and Violante (2010) show that BPP is highly robust at recovering the true MPC to temporary shocks parameter in a variety of models. However, they do not explore models with housing, so one might wonder whether BPP recovers the true MPC to temporary income shocks in our environment. We have performed exercises similar to Kaplan and Violante (2010) in our model and indeed found that this procedure continues to work well even in the presence of housing with transaction costs and rental markets.

⁴⁹Abowd and Card (1989) show that this parsimonious specification fits income data well.

⁵⁰That is, the BPP estimator gives an estimate of elasticities, but we maintain their notation to avoid confusion.

⁵¹We then convert this MPC in logs to the MPC in levels relevant for our theory by multiplying by the average value of C/Y in each bin. This also means that the elasticity of consumption to house prices shown below is equal to the MPC estimate from BPP times the average of $\frac{PH}{Y}$ in each bin.

⁵²An alternative approach to estimate MPCs, used by Johnson et al. (2006), uses random government rebate timing and CEX data. Unfortunately, it is well known that the resulting standard errors on MPCs are large since CEX data has smaller sample sizes than PSID, no panel structure to disentangle true changes in consumption from measurement error, and covers a smaller fraction of consumption. Large standard errors are especially problematic for us, since we need to estimate MPCs conditional on different levels of housing wealth.

4.2 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). This is why we use 1999 as the first year of our sample.

Since we use almost the same underlying sample as Kaplan et al. (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 25-60 years old. Our final sample has 30,462 observations over the pooled years 1999-2011 (seven sample years).

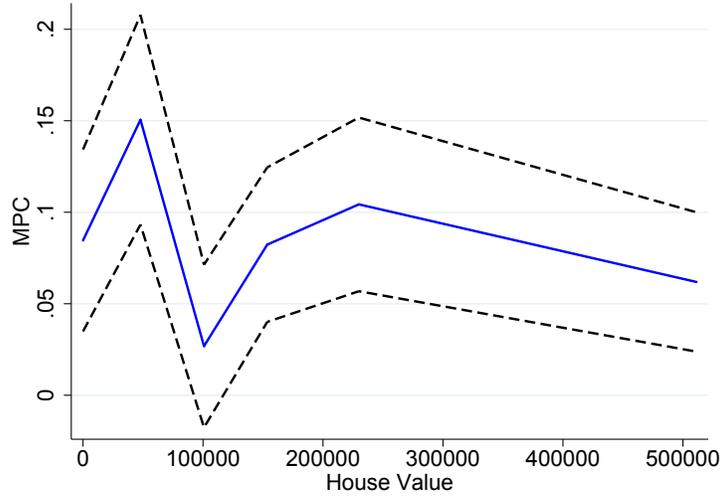
We use the same consumption definition as Blundell et al. (2014) including food at and away from home, utilities, gasoline, car maintenance, public transportation, childcare, health expenditures and education. This data covers approximately 70% of consumption in the CEX. We define income as the sum of labor income and government transfers. We purge the data of non-model features by regressing $\ln c_{it}$ and $\ln y_{it}$ on year and cohort dummies, education, race, family structure, employment, geographic variables, and interactions of year dummies with education, race, employment, and region.

4.3 Results

Since our formula is equal to the product of the MPC and the value of housing, we first examine how the MPC varies with the value of housing. To do this, we estimate the MPC using the BPP methodology separately in 6 housing bins. The first bin includes only renters—that is, households who own zero housing. The other five bins are constructed as the quintiles of the home value distribution and so are of equal size by construction. In total, in our database we have approximately 11,000 renters and 19,000 home owners.

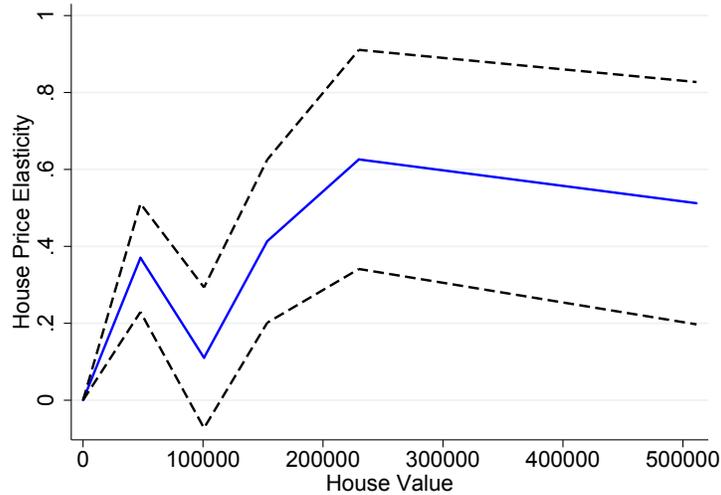
Figure 12 shows the relationship between the MPC and home values. Overall the relationship between home values and the MPC is flat or slightly declining (though this decline is not statisti-

Figure 12: MPC and Housing



cally significant). Therefore, housing-rich households in the data still display high MPCs, which will contribute to a large aggregate elasticity. Figure 13 translates these results into elasticities by bins, which are computed by multiplying the MPC by bin in Figure 12 times the ratio of average home values over average consumption in the same bin. Elasticities for homeowners are non-monotone in housing and range between 0.11 and 0.63.

Figure 13: Elasticity and Housing



Using these estimates we can compute the aggregate elasticity of non-durable consumption for the whole sample, which is equal to 0.33.⁵³ While our earlier modeling results emphasized heterogeneity in responses across households at a given point in time, the distribution of $MPCs$

⁵³This value is given by $(\sum_k MPC_k \times PH_k)/C$, where MPC_k is the estimated MPC in bin k , PH_k is the average home value for bin k , and C is aggregate consumption.

and housing can also potentially move across time. For example, house prices and leverage increased substantially over the housing boom, with potential effects on the aggregate elasticity of consumption to house prices. Indeed, we find that the overall average elasticity masks significant time-series variation: the elasticity is 0.31 during the housing boom (2002-2006) and 0.45 during the housing bust (2008-2010).⁵⁴ We will show in the next section that our model naturally generates larger consumption responses when housing booms change to busts.

This sufficient-statistic based measure of housing price effects provides an alternative to the standard empirical strategies for estimating such effects. Existing evidence of the importance of house price movements for consumption relies either on strong time-series identification assumptions or on the identification of “exogenous” changes in house prices at the micro level. Amongst the latter style, the series of papers by Mian and Sufi using the elasticity of housing supply to instrument for local house price movements has probably been the most influential in arguing for the causal effects of house prices on consumption. However, some recent papers such as Davidoff (2013) have questioned the validity of supply elasticity in this context and by extension the evidence for causal housing price effects. Our sufficient-statistic based measure of housing price effects relies on a completely different source of identification, so it does not require us to take any stance on the validity of this IV strategy. Nevertheless, despite our alternative identification approach, we ultimately find consumption elasticities to house price changes that are large and similar to those in Mian et al. (2013).⁵⁵

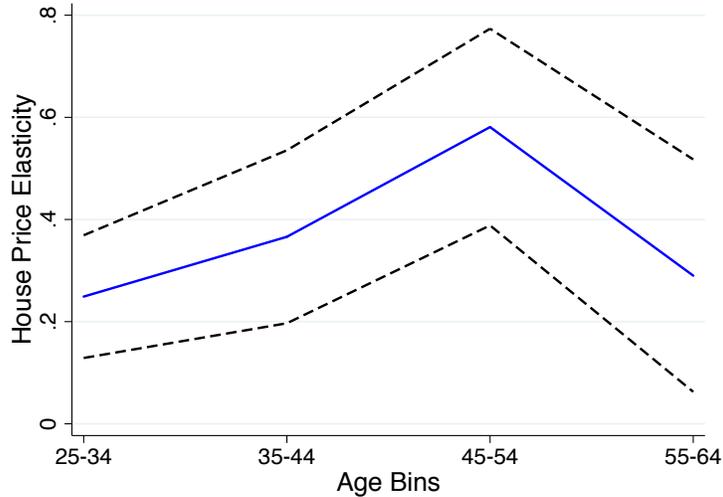
To be clear, this sufficient-statistic approach has pros and cons relative to the existing literature. While we show that it is a valid approach in many commonly used modeling environments, it does rely on the general validity of these models. It also requires estimates of the MPC out of temporary income shocks, which are themselves somewhat controversial. However, we avoid the need for exogenous variation in house prices, which has been a source of frequent debate in the existing literature. Although the focus of our paper is not empirical, we nevertheless find it reassuring that these complementary approaches ultimately deliver quite similar results, which are also consistent with the most realistic calibrations of our theoretical model.

Next, we explore the relationship between the non-durable consumption elasticity and the life cycle. Ideally, we would like to estimate an elasticity for each age. However, due to the limited sample size, we group agents into the following four age bins: 25-34, 35-44, 45-54 and 55-64. The

⁵⁴Since the increase in housing values in the boom is roughly undone by the decrease in housing values in the bust, average housing values are similar in both periods and the difference in elasticities is mostly driven by greater MPCs during the bust. This is consistent with the model predictions shown below.

⁵⁵Although it is important to note the caveat from footnote 14, that the Mian et al. (2013) measures include local general equilibrium effects. Notice also that our theoretical result validates the housing wealth measure in Mian et al. (2013). In that paper, they measure local changes in housing net worth as $\Delta \log_{06-09}^{H,i} * H_{2006}^i / NW_{2006}^i$, where H_{2006}^i is the value of the housing stock owned by households in zip code i in 2006. That is, their housing net worth shock multiplies the change in house prices by the gross value of housing rather than by housing equity. One might intuitively think that the strength of housing price effects for a given house price change might depend on housing equity rather than gross housing, but our simple formula shows that this intuition is incorrect.

Figure 14: Elasticity and Age



results are reported in Figure 14, which shows a pronounced hump shape in the relation between age and the elasticity of consumption to house prices. This hump shape reflects two separate forces. First, homeownership rates and housing holdings are hump shaped in age. Second, MPCs are slightly hump shaped in age.

Recall, from Figure 8, that the model also displays a hump-shaped relation between elasticity and age. The size of the elasticities produced by the model is fairly close to the data although the hump shape is somewhat more pronounced in the data.⁵⁶

5 A Boom-Bust Episode

So far we have focused on the partial equilibrium response of consumption to a one-time exogenous change in house prices. In this section, we extend the analysis to a simple dynamic general equilibrium exercise. The main objective of the exercise is to show that the large partial equilibrium responses obtained so far in our baseline model with fixed costs and a rental option can also lead to large general equilibrium responses.⁵⁷ An additional objective of the exercise is to show that the response of an economy to a change in house prices depends on the history of past shocks. In particular, we show that an economy coming from a boom in house prices is more sensitive to a following reduction in prices.

⁵⁶The model does not generate enough middle-aged consumers with high housing values and high MPCs, or, in the parlance of Kaplan and Violante (2014), enough wealthy-hand-to-mouth consumers. This occurs in part because while Kaplan and Violante (2014) rule out borrowing against the illiquid asset, we allow for such borrowing, subject to a refinancing cost calibrated to match cash-out shares in the data. It is an interesting open question how to produce wealthy-hand-to-mouth behavior while allowing for realistic access to home equity financing.

⁵⁷Naturally, this conclusion depends on the nature of the shock and on the way in which the general equilibrium model is specified. Kaplan et al. (2015) make this point clear by looking at the effects of a variety of shocks in a related model. See the introduction for additional discussion of the relationship between our results.

To determine house prices in general equilibrium, we combine the model developed so far with a competitive construction sector which transforms consumption goods into houses. Given that our focus is on the demand side, we use a very stylized model of housing supply. The construction technology is described by the quadratic cost function $G(I_t, \zeta_t) = \zeta_t I_t + (\kappa/2)I_t^2$, where $G(I_t, \zeta_t)$ are the consumption goods needed as inputs to build I_t new housing units and ζ_t is a supply shock. Optimization by firms in the construction sector implies $P_t = \kappa I_t + \zeta_t$. Market clearing in the housing market requires $I_t = H_t - (1 - \delta)H_{t-1}$, where H_t is the aggregate demand for housing by households.⁵⁸

On the demand side, we introduce a shock to agents' expectations, capturing a form of irrational bubble. We assume that the economy starts in steady state with no shocks. At date $t = 0$, agents become optimistic about future house price growth, this leads to an increase in housing demand and thus to an actual increase in house prices today. Agents remain optimistic for a number of periods and eventually revert to zero expected house price appreciation.⁵⁹ Expectations are modeled as follows. In each period t , households observe the current price P_t and form deterministic expectations of future prices:

$$E_t[P_{t+j}] = P_t e^{g_t \frac{1-\lambda^j}{1-\lambda}} \quad \text{for } j = 1, 2, \dots \quad (5)$$

Households expect prices to grow at the rate g_t in the next period and to keep growing gradually towards the long-run level $P_t e^{g_t/(1-\lambda)}$. Expectations are not rational here, as they do not match the realized equilibrium path.

Given a sequence $\{g_t, P_t\}$, the optimization problem of the household can be solved at each t , using (5) to form expectations. An equilibrium is then given by a sequence of prices P_t and a sequence of joint distributions of housing and non-housing wealth, such that households optimize, construction firms optimize, the housing market clears in each period, and the distribution of housing and non-housing wealth at $t + 1$ is given by the distribution at t and by the agents' optimal policies.⁶⁰

The parameters for the households' side are exactly as in the full model with rentals and housing transaction costs of Section 3.⁶¹ The parameter κ in the supply equation is set at 36.9.⁶²

⁵⁸Here we are only endogenizing house prices and not the interest rate or household incomes Y_{it} . We can think of our general equilibrium economy as a small open economy that takes the world interest rate r as given and where households' incomes are unaffected by domestic spending. Of course, a full understanding of the relation between house prices and real activity would require a richer general equilibrium model that also endogenizes interest rates and incomes.

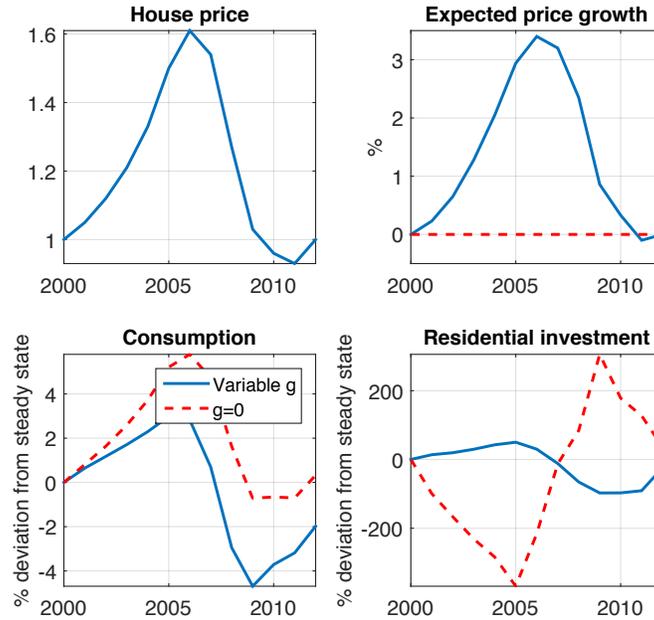
⁵⁹Changes in expected house appreciation are consistent with recent empirical work by Case et al. (2012), Burnside et al. (2016) and Bailey et al. (2016).

⁶⁰This definition of equilibrium without rational expectations is analogous to the temporary general equilibrium notion of Grandmont (1977), which was recently revived by Piazzesi and Schneider (2008).

⁶¹In Appendix 5 we do the same exercise in the frictionless model and show that it delivers similar but quantitatively larger effects.

⁶²The number comes from a simple linear regression of I_t on P_t in 2000-2005.

Figure 15: Boom-Bust Simulation



The values for the demand and supply shocks g_t, ζ_t are chosen to match the dynamics of house prices and aggregate residential investment in the US in 2000-2013. For house prices, we use the average US home price series from Shiller (2015).⁶³ The residential investment rate is taken from NIPA.⁶⁴ The value of λ in the expectation equation (5) is set at 0.5.

Figure 15 reports the results of the exercise (solid blue lines). For comparison with the previous sections, we also compute responses with no expected appreciation, that is, we let prices follow the path in the top-left panel and assume consumers perceive each change in P_t as unexpected and permanent, so $g_t = 0$ for all t . The dashed red lines correspond to this case.

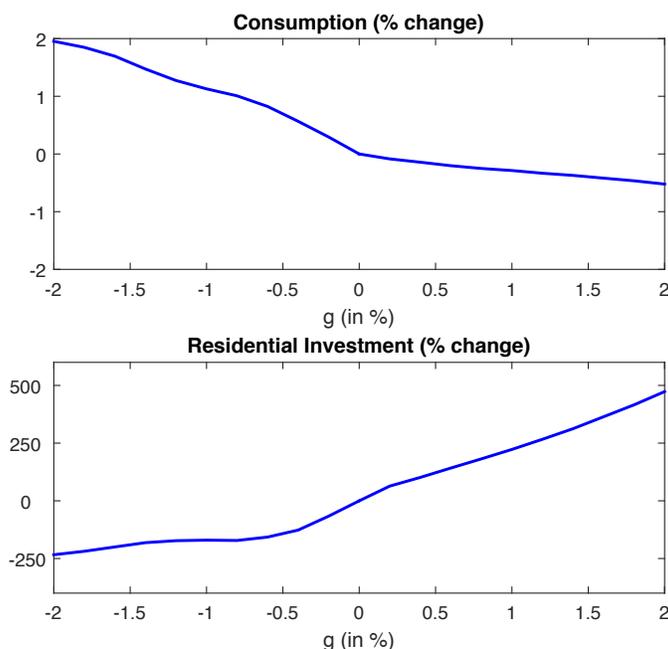
Figure 15 leads us to two main observations: the introduction of the g_t shocks leads to very different predictions regarding housing demand, but to relatively similar predictions regarding consumption.

The first observation comes from comparing the solid blue and the dashed red lines in the bottom-right panel. When $g_t = 0$ the model predictions for housing demand are completely counterfactual, with the demand for residential investment falling in the boom and increasing in the bust. To produce a boom-bust in residential investment, the model asks for a demand shock on the household side. The expectation shock g_t provides such a shock.

⁶³We use the Real Home Price Index from www.econ.yale.edu/shiller/data, file Fig2-1.xls.

⁶⁴In the model, the residential investment rate is $H_{t+1}/H_t - (1 - \delta)$. In the data, it is measured by residential investment (BEA Fixed Asset Table 1.5, line 8) divided by the stock of private residential capital (BEA Fixed Assets Table 1.1, line 14). Since the model has no long-run growth, the steady-state investment rate is $\delta = 0.022$, which is lower than the data average of 0.046 (in the period 1960-2013). We adjust for unmodeled growth by subtracting 0.024 from the investment rate in the data.

Figure 16: Effects of Expected House Price Appreciation



What are the consumption implications of a *combined* boom in house prices together with a positive shock to housing demand? The answer to this question is in the bottom-left panel of Figure 15, which shows the behavior of consumption over the house price boom-bust with and without g shocks. Overall, shocks to g result in a modest dampening of the consumption boom followed by an amplification of the consumption bust, so that consumption falls substantially further below steady-state as prices return to normal than in the economy with no housing demand shocks. Notice also that the difference between the two economies grows over the boom-bust period. Overall, this reflects the fact that g has little direct effect on consumption and mainly works indirectly by changing the endogenous distribution of wealth in the economy. Importantly, this small direct effect of g means that our sufficient-statistic remains useful even in general equilibrium, since it implies that consumption is mostly determined by the partial equilibrium house price changes analyzed in the rest of the paper.

Let us provide some intuition for why the introduction of the g_t shock has large direct effects on housing demand but small direct effects on consumption. Figure 16 shows what happens on impact to an economy in steady-state hit by an unexpected, one-time change in g . The top panel shows the response of consumption, the bottom panel shows the response of residential investment. The figure shows that an increase in g tends to mildly reduce consumption for given initial wealth, while it dramatically increases residential investment. The intuition is as follows: an increase in g implies an expected increase in the cost of housing services in the future and thus an expected increase in the overall CPI. This has both an intertemporal substitution effect, which

increases current consumption, and an income effect, which lowers it. When $\sigma > 1$ the income effect dominates and consumption falls. However, as long as σ is not too large, the net effect is small. The effect of g on spending on housing services is of the same sign and magnitude as the effect on non-durable consumption, due to Cobb-Douglas preferences. Spending on housing services is, in terms of implicit rental costs,

$$\left[1 - e^g \frac{1 - \delta}{1 + r}\right] PH.$$

So while this expression overall is decreasing slightly, the level of H increases since the expression in square brackets is strongly decreasing in g . This comparative statics exercise helps explain why in Figure 15, the introduction of the g_t shocks has large effects on residential investment and modest effects on consumption.

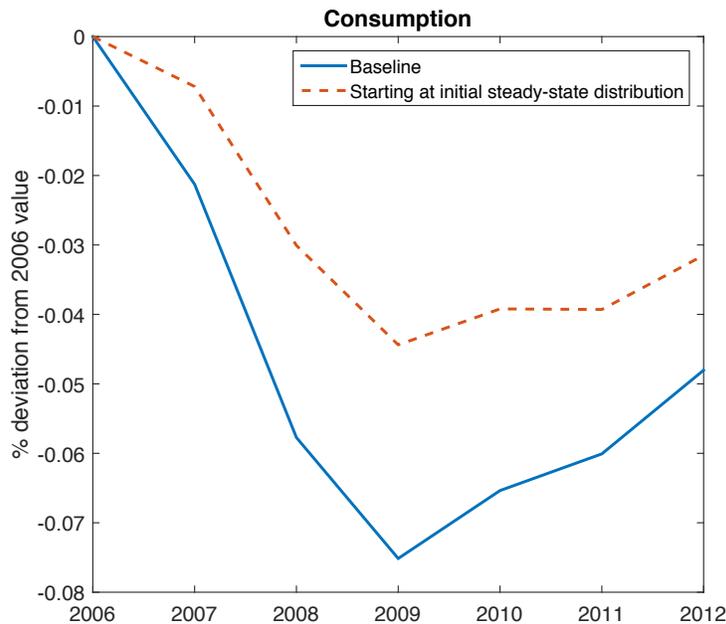
To further understand the consumption responses in Figure 15, we make two additional remarks on the dynamics.

First, notice that in the previous sections, we have focused on the impact effect of a house price change on consumption. However, our model predicts that the effect of a permanent house price change on consumption is *transitory*: as households converge back to their target wealth levels, consumption gradually reverts to its steady-state level, which is independent of P . Therefore, when we look at the cumulative effect of a sequence of house price increases, such as the one between 2000 and 2006, the total increase in consumption is smaller than the sum of the impact effects.⁶⁵ The drop in consumption after 2006 is larger than the prior increase (and is roughly in line with the decline in consumption observed in the recession). This occurs for three reasons: first, absent price changes, as we just argued, consumption would already be on a downward path reverting to its steady-state; second, a price drop actually takes place and is faster than the increase; third, the reduction in housing wealth, for given debt levels, leads to larger MPCs. The third effect is especially interesting and we discuss it now in more detail.

In a dynamic equilibrium the endogenous change in the distribution of state variables A_{it}, H_{it} means that the same current shock will have different effects depending on the history of past shocks. Figure 17 compares the consumption response in the bust period in our exercise (solid blue) to a counterfactual response if we assume that the distribution of households' net worth at the beginning of the bust is unchanged from the initial steady-state (dashed red line). The figure confirms that the endogenous change in the wealth distribution over the boom has a quantitatively important effect in the bust, leading to a deeper and more prolonged contraction in consumption. This occurs because the boom leads to an accumulation of housing wealth and of debt, as shown in Figure 18. This simultaneously increases MPC and H and thus makes household balance sheets

⁶⁵This explains why, even in the partial equilibrium exercise, a total price increase of 60% yields a total increase in consumption of about 5%, which is smaller than the predictions one would get from the impact elasticity. Using the impact elasticity in the model with rentals and adjustment costs would give $14.4\% = 0.24 \times 60\%$.

Figure 17: Response to a House Price Drop: The Effect of the Initial Distribution



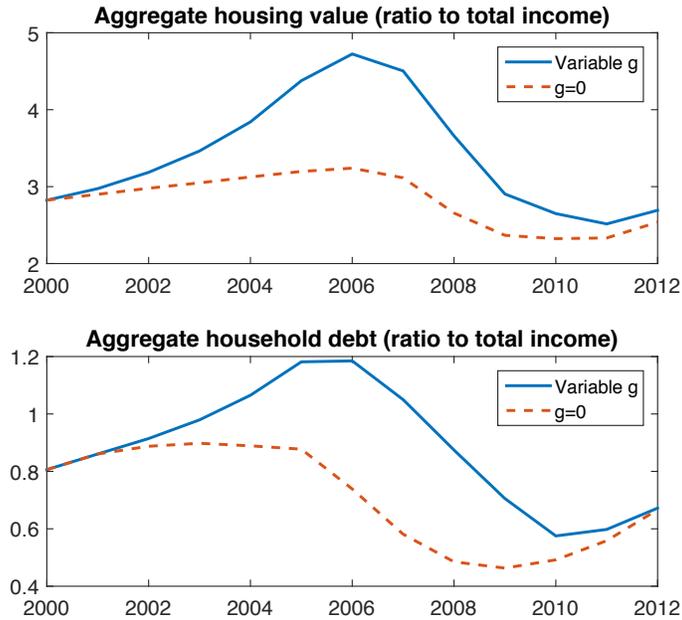
more sensitive to a drop in house prices through the mechanism explored in previous sections. Notice that the presence of the g_t shocks amplifies the importance of these effects, since increases in g_t lead to greater accumulation of housing and debt in the boom phase.

Thus, a useful lesson from this exercise is that endogenous time-variation in the distribution of debt leads to changes in the strength of house price effects on consumption. This implies that accounting for variation in micro level heterogeneity across time is crucial for making accurate predictions about the aggregate effects of housing market shocks. This can help rationalize why studies using data from different time periods often arrive at substantially different conclusions as to the importance of housing price effects. For example, Case et al. (2013) finds substantially larger housing price effects when using data from 1975-2012 than in their otherwise identical previous study, Case et al. (2005), which only included data through 1999.

The overall take-out from the boom-bust exercise is as follows: Introducing shocks to expected appreciation can deliver a boom-bust in residential investment together with a boom-bust in house prices. The direct effect of shocks to expected appreciation is to mildly dampen the boom-bust in consumption through intertemporal channels. However, the same shocks have a significant indirect effect, which tends to amplify the boom-bust through endogenous changes in the distribution of housing and debt. Since the direct effect is small, our earlier formula remains useful.

The analysis in this section also indirectly sheds light on the effect of transitory shocks to house prices in our model. Transitory shocks can be interpreted as the combination of a one-time

Figure 18: Boom-Bust Simulation: Housing Wealth and Debt



increase in P combined with a negative g shock. In our calibration, we just saw that a negative g shock mildly increases current consumption, so a temporary house price shock will have a larger impact effect on consumption than a permanent shock. On the other hand, the effect will go away more quickly, when house prices start to fall in the periods following the initial increase.

6 Conclusion

In this paper, we explore the implications of consumption theory for understanding housing price effects on consumption. While a large and growing empirical literature documents strong responses of consumption to identified house price movements, a large theoretical literature argues that this response should be small.

We show that a relatively standard incomplete markets consumption model with idiosyncratic risk actually implies quite large elasticities, and we derive a new sufficient statistic to provide intuition for this large value. This sufficient statistic continues to work well in more realistic models, and we show that these models deliver elasticities closely in line with empirical evidence. We then take our sufficient statistic directly to micro data and show it yields similar conclusions. Not only is our sufficient statistic useful for understanding why elasticities in our model are large, it also provides good intuition for why some previous models have found small elasticities. These models often either imply counterfactually small MPCs or shut down some of the theoretical channels through which house prices can affect consumption.

Our results also imply that the size of housing price effects may vary substantially across households and time. Indeed, we find large variation over the boom-bust in our model and find that elasticities vary substantially with age and home ownership status. The presence of this time-variation and heterogeneity may help reconcile various empirical results and should be accounted for when predicting the aggregate consequences of housing market policies.

References

- Abowd, John M. and David Card**, “On the Covariance Structure of Earnings and Hours Changes,” *Econometrica*, 1989, 57 (2), 411–445.
- Aguiar, Mark and Erik Hurst**, “Deconstructing Lifecycle Expenditures,” *Journal of Political Economy*, 2013, 212 (1).
- Attanasio, Orazio and Guglielmo Weber**, “The UK consumption boom of the late 1980s: aggregate implications of microeconomic evidence,” *The Economic Journal*, 1994, pp. 1269–1302.
- , **Andrew Leicester, and Matthew Wakefield**, “Do House Prices Drive Consumption Growth? The Coincident Cycles of House Prices and Consumption in the UK,” *Journal of the European Economic Association*, 2011, 9 (3).
- , **Laura Blow, Robert Hamilton, and Andrew Leicester**, “Booms and Busts: Consumption, House Prices, and Expectations,” *Economica*, 2009, 76 (301), 20–50.
- Auclert, Adrien**, “Monetary Policy and the Redistribution Channel,” 2015. mimeo, Stanford University.
- Bailey, Mike, Ruiqing Cao, Theresa Kuchler, and Jonannes Stroebel**, “Social Networks and Housing Markets,” Working Paper 22258, National Bureau of Economic Research 2016.
- Berger, David and Joseph Vavra**, “Consumption Dynamics During Recessions,” *Econometrica*, 2015, 83 (1), 101–154.
- Bhutta, Neil and Benjamin J. Keys**, “Interest Rates and Equity Extraction During the Housing Boom,” *University of Chicago Kreisman Working Papers Series in Housing Law and Policy*, 2014, 3.
- Blundell, Richard, Luigi Pistaferri, and Ian Preston**, “Consumption Inequality and Partial Insurance,” *American Economic Review*, 2008, 98 (5), 1887–1921.
- , – , and **Itay Saporta-Eksten**, “Consumption Inequality and Family Labor Supply,” *2014 Meeting Papers No. 322 Society for Economic Dynamics*, 2014.
- Buiter, Willem H.**, “Housing Wealth Isn’t Wealth,” *NBER Working Paper 14204*, 2008.
- Burnside, Craig, Martin Eichenbaum, and Sergio Rebelo**, “Understanding Booms and Busts in Housing Markets,” *Journal of Political Economy*, 2016, pp. 1088–1147.

- Caballero, Ricardo and Eduardo M. R. A. Engel**, “Explaining Investment Dynamics in U.S. Manufacturing: A Generalized (S,s) Approach,” *Econometrica*, 1999, *67* (4), 783—826.
- Campbell, John Y. and Joao F. Cocco**, “How Do House Prices Affect Consumption?,” *Journal of Monetary Economics*, 2007, *54* (3), 591—621.
- Carroll, Chris and Wendy Dunn**, “Unemployment Expectations, Jumping (S,s) Triggers, and Household Balance Sheets,” *NBER Macroeconomics Annual*, 1998, *12*, 165—230.
- Carroll, Christopher D., Misuzu Otsuka, and Jiri Slacalek**, “How Large Are Housing and Financial Wealth Effects? A New Approach,” *Journal of Money Credit and Banking*, 2011, *1* (43), 55—79.
- Case, Karl E., John M. Quigley, and Robert J. Shiller**, “Comparing Wealth Effects: The Stock Market versus the Housing Market,” *Advances in Macroeconomics*, 2005, *5* (1).
- , —, and —, “Wealth Effects Revisited: 1975-2012,” *Critical Finance Review*, 2013, *2* (1), 101—128.
- , **Robert J. Shiller, and Anne Thompson**, “What Have They Been Thinking? Home Buyer Behavior in Hot and Cold Markets,” *NBER Working Paper 18400*, 2012.
- Cerutti, Eugenio, Jihad Dagher, and Mr Giovanni Dell’Ariccia**, “Housing Finance and Real-Estate Booms: A Cross-Country Perspective,” Staff Discussion Notes 15/12, International Monetary Fund 2015.
- Chen, Hui, Michael Michaux, and Nikolai Roussanov**, “Houses as ATMs? Mortgage Refinancing and Macroeconomic Uncertainty,” Working Paper 19421, National Bureau of Economic Research 2013.
- Chetty, Raj**, “Sufficient Statistics for Welfare Analysis: A Bridge Between Structural and Reduced-Form Methods,” *Annual Review of Economics*, 2009, *1*, 451—488.
- Cocco, Joao**, “Portfolio Choice in the Presence of Housing,” *The Review of Financial Studies*, 2005.
- Davidoff, Thomas**, “Supply Elasticity and the Housing Cycle of the 2000s,” *Real Estate Economics*, 2013, *41* (4), 793—813.
- Davis, Morris A. and François Ortalo-Magné**, “Household Expenditures, Wages, Rents,” *Review of Economic Dynamics*, 2011, *14*, 248—261.
- DeFusco, Anthony**, “Homeowner Borrowing and Housing Collateral: New Evidence from Expiring Price Controls,” 2015.

- Díaz, Antonia and María José Luengo-Prado**, “The Wealth Distribution with Durable Goods,” *International Economic Review*, 2010, 51 (1), 143—170.
- Eggertsson, Gauti B and Paul Krugman**, “Debt, deleveraging, and the liquidity trap: A fisher-minsky-koo approach*,” *The Quarterly Journal of Economics*, 2012.
- Favilukis, Jack, Sydney C. Ludvigson, and Stijn Van Nieuwerburgh**, “The Macroeconomic Effects of Housing Wealth, Housing Finance, and Limited Risk-Sharing in General Equilibrium,” Working Paper, NYU 2015.
- Floden, Martin and Jesper Lindé**, “Idiosyncratic Risk in the United States and Sweden: Is There a Role for Government Insurance,” *Review of Economic Dynamics*, 2001, 4 (2), 406—437.
- Glaeser, Edward L.**, “Real Estate and The Macroeconomy: Discussion,” *Brookings Papers on Economic Activity*, 2000.
- Gorea, Denis and Virgiliu Midrigan**, “Liquidity Constraints in the U.S. Housing Market,” 2015.
- Grandmont, Jean Michel**, “Temporary General Equilibrium Theory,” *Econometrica*, 1977, 45 (3), 535–572.
- Guerrieri, Luca and Matteo Iacoviello**, “Collateral Constraints and Macroeconomic Asymmetries,” *Manuscript, Federal Reserve Board*, 2014.
- Guerrieri, Veronica and Guido Lorenzoni**, “Credit Crises, Precautionary Savings, and the Liquidity Trap,” Working Paper, National Bureau of Economic Research 2011.
- Guiso, Luigi, Monica Paiella, and Ignazio Visco**, “Do capital gains affect consumption? Estimates of wealth effects from italian households’ behavior,” Temi di discussione 555, Bank of Italy 2005.
- Guvenen, Fatih and Anthony A. Smith**, “Inferring Labor Income Risk and Partial Insurance from Economic Choices,” *Econometrica*, 2014, 82 (6), 2085—2129.
- Hall, Robert E**, “The Long Slump,” *American Economic Review*, 2011, 101 (2), 431–69.
- Huo, Zhen and José-Víctor Ríos-Rull**, “Balance sheet recessions,” Technical Report, Working Paper, Federal Reserve Bank of Minneapolis 2013.
- Johnson, David S., Jonathan A. Parker, and Nicholas S. Souleles**, “Household Expenditure and the Income Tax Rebates of 2001,” *American Economic Review*, 2006, 96 (5), 1589—1610.

- Justiniano, Alejandro, Giorgio E Primiceri, and Andrea Tambalotti**, “Household leveraging and deleveraging,” *Review of Economic Dynamics*, 2015, 18 (1), 3–20.
- Kaplan, Greg and Gianluca Violante**, “How Much Consumption Insurance Beyond Self-Insurance,” *American Economic Journals: Macroeconomics*, 2010, 2 (4), 53–87.
- and – , “A Model of the Consumption Response to Fiscal Stimulus Payments,” *Econometrica*, 2014, 82 (4), 1199–1239.
- , **Giovanni L. Violante, and Justin Weidner**, “The Wealthy Hand-To-Mouth,” *NBER Working Paper 20073*, 2014.
- , **Kurt Mitman, and Gianluca Violante**, “Consumption and House Prices in The Great Recession,” 2015.
- Mian, Atif, Kamalesh Rao, and Amir Sufi**, “Household Balance Sheets, Consumption, and the Economic Slump,” *The Quarterly Journal of Economics*, 2013, 128 (4), 1687–1726.
- Midrigan, Virgiliu and Thomas Philippon**, “Household Leverage and the Recession,” Technical Report 16965, NBER 2011.
- Petev, Ivaylo, Luigi Pistaferri, and Itay Saporta Eksten**, “Consumption and the Great Recession: An analysis of trends, perceptions, and distributional effects,” in D. Grusky, B. Western, and C. Wimer, eds., *Analyses of the Great Recession*, 2012.
- Piazzesi, Monika and Martin Schneider**, “Inflation Illusion, Credit and Asset Pricing,” in John Y. Campbell, ed., *Asset Pricing and Monetary Policy*, Chicago University Press, 2008, pp. 147–181.
- , – , and **Selale Tuzel**, “Housing, Consumption and Asset Pricing,” *Journal of Financial Economics*, 2007, 83, 531–569.
- Saiz, Albert**, “The Geographic Determinants of Housing Supply,” *The Quarterly Journal of Economics*, 2010, 125 (3), 1253–1296.
- Shiller, Robert J.**, “Irrational Exuberance,” *Princeton University Press*, 2015.
- Sinai, Todd and Nicholas S Souleles**, “Owner-Occupied Housing as a Hedge against Rent Risk,” *The Quarterly Journal of Economics*, 2005, pp. 763–789.
- Stroebel, Johannes and Joseph Vavra**, “House Prices, Local Demand, and Retail Prices,” 2015.

- Tauchen, George**, “Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions,” *Economics Letters*, 1986, *20* (2), 177—181.
- Vavra, Joseph**, “Inflation Dynamics and Time-Varying Volatility: New Evidence and an Ss Interpretation,” *The Quarterly Journal of Economics*, 2014, *129* (1), 215—258.
- Winberry, Thomas**, “Lumpy Investment, Business Cycles, and Stimulus Policy,” 2015.
- Yao, Rui and Harold H. Zhang**, “Optimal Consumption and Portfolio Choices with Risky Housing and Borrowing Constraints,” *The Review of Financial Studies*, 2005.

Appendix

HOUSE PRICES AND CONSUMER SPENDING

D. Berger, V. Guerrieri, G. Lorenzoni, J. Vavra

1 Decomposition

In this section we derive decompositions of the total house price effect into income, substitution, collateral and endowment effects and discuss their interpretation both for the simple PIH model of Subection 2.2 and for the general baseline model treated in the rest of Section 2.

1.1 PIH model

For simplicity, we focus on the infinite-horizon version of the model and drop the individual subscript i . The household's utility function is $\sum_{t=0}^{\infty} \beta^t U(C_t, H_t)$. The per-period budget constraints, with a no Ponzi condition, can be aggregated into the intertemporal budget constraint at date 0:

$$\sum_{t=0}^{\infty} q^t [C_t + [P_t - q(1 - \delta) P_{t+1}] H_t - Y_t] = (1 + r) A_{-1} + (1 - \delta) P_0 H_{-1},$$

where $q \equiv 1/(1 + r)$. Define the implicit rental rate

$$R_t \equiv P_t - q(1 - \delta) P_{t+1}.$$

Suppose that the house price is initially constant at P so that the implicit rental rate is then $R = (1 - q(1 - \delta)) P$. We then want to decompose the response of C_0 to a permanent change in P .

Define the Marshallian demand for C_0

$$C_0 = C_{0,m}(R, I),$$

which comes from maximizing $\sum \beta^t U(C_t, H_t)$ subject to $\sum q^t [C_t + RH_t] = I$. Similarly, define the Hicksian demand

$$C_0 = C_{0,h}(R, U),$$

which comes from minimizing $\sum q^t [C_t + RH_t]$ subject to $\sum \beta^t U(C_t, H_t) \geq U$. The effect of a permanent change dR is equivalent to summing the effects of identical changes in the price of housing services dR_t . So the standard decomposition result also applies and implies that the total

response of consumption to a change in P satisfies

$$\frac{dC_0}{dP} = [1 - q(1 - \delta)] \frac{\partial C_{0,h}(R, U)}{\partial R} - [1 - q(1 - \delta)] \frac{\partial C_{0,m}(R, I)}{\partial I} \sum q^t H_t + \frac{\partial C_{0,m}(R, I)}{\partial I} (1 - \delta) H_{-1}.$$

The first term is the substitution effect, the second the income effect, the third the endowment effect.

Under the assumption of Cobb-Douglas preferences and $q = \beta$, the solution to the household problem gives a constant level of housing $H_t = H_0$ and of consumption $C_t = C_0$ with

$$C_0 = \alpha \frac{r}{1+r} \left[(1+r) A_{-1} + (1-\delta) P_0 H_{-1} + \sum q^t Y_t \right].$$

Now consider a household that starts at $H_{-1} = H_0$ (which applies for small changes in R , for a household that started life at any time $t < 0$). Then the sum of the income and endowment effects is

$$\frac{\partial C_{0,m}(R, I)}{\partial I} \left\{ -[1 - q(1 - \delta)] \sum q^t + (1 - \delta) \right\} H = -\frac{\delta}{1-q} \frac{\partial C_{0,m}(R, I)}{\partial I} H < 0.$$

The net effect is small if δ is small, and exactly zero if $\delta = 0$. This captures the Sinai and Souleles (2005) intuition that when the price of houses increases, and the net present value of implicit rental cost increases by the same amount, then there are small wealth effects. That the net effect is actually negative reflects the fact that the consumer is a net buyer of housing since he needs to replace the depreciated fraction of housing in all periods.

At the same time, we know that

$$\frac{dC_0}{dP} = \alpha \frac{r}{1+r} (1 - \delta) H_{-1} = \frac{\partial C_{0,m}(R, I)}{\partial I} (1 - \delta) H_{-1} > 0.$$

This implies that the total effect can be interpreted in two ways:

- a substitution effect that more than compensates for the negative net of income and endowment effects derived above;
- a pure endowment effect, with the substitution and income effects canceling each other.

1.2 General model with no adjustment costs

Given the household's optimization problem on page 10, with $P_t = P$ for all t , let $C(W, s; P, \theta)$ denote the optimal consumption policy. Consider the effects of an unexpected, permanent change in price from P_0 to P_1 . Let (H_-, A_-) denote the household initial holdings of housing and of the

risk-free asset. Household net wealth is then

$$W_0 = P_0 (1 - \delta) H_- + (1 + r) A_-$$

before the shock, and

$$W_1 = P_1 (1 - \delta) H_- + (1 + r) A_-$$

after the shock. The total consumption change is

$$\Delta C \equiv C(W_1, s; P_1, \theta) - C(W_0, s; P_0, \theta).$$

To compute the collateral effect choose $\hat{\theta}$ such that

$$(1 - \theta) P_1 = (1 - \hat{\theta}) P_0,$$

that is consider a change in the collateral requirement that exactly offsets the change in price. The collateral effect is then defined as

$$CE \equiv C(W_0, s; P_0, \hat{\theta}) - C(W_0, s; P_0, \theta).$$

Next, find the value of initial wealth \hat{W} such that $V(\hat{W}, s; P_1, \theta) = V(W_0, s; P_0, \hat{\theta})$. That is, we find the wealth that keeps utility unchanged after a change in the price of housing. Here, we are adapting the logic of Hicksian compensation to our dynamic, incomplete-markets problem. We define the substitution effect as

$$SE \equiv C(\hat{W}, s; P_1, \theta) - C(W_0, s; P_0, \hat{\theta}),$$

where the change in θ is introduced to mute the collateral effect and the change in W is introduced to mute the income effect. The income effect is given by

$$IE \equiv C(W_0, s; P_1, \theta) - C(\hat{W}, s; P_1, \theta).$$

Finally, the endowment effect is just given by

$$EE \equiv C(W_1, s; P_1, \theta) - C(W_0, s; P_1, \theta),$$

and elementary algebra gives

$$\Delta C = SE + IE + CE + EE,$$

and since $C(W, s; P, \theta)$ is independent of P , as shown in Proposition 1, we have $\Delta C = EE$. Notice that the exact result of the decomposition, for discrete changes in P , depends on the ordering (e.g., here we started from the collateral effect).

2 Extension: CES Preferences

In this section, we extend the analysis to CES preferences for consumption and housing. In the body of the paper, we assume Cobb-Douglas utility—i.e., elasticity of substitution equal to 1—and use that assumption to derive Proposition 1. Here we show that the proposition can be extended to the case of CES preferences if we make the additional assumption of $\theta = 0$, that is, if we consider a very loose collateral requirement. In that case, our sufficient statistic formula extends naturally by adding a new term that is positive in the case of elasticity of substitution bigger than 1 and negative in the opposite case. We then provide numerical results for the case $\theta > 0$, calibrated as in the main text, and show that, for reasonable values of the elasticity of substitution between consumption and housing, our amended sufficient-statistic formula provides a good approximation and the additional term has small effects if the elasticity of substitution is in the range 0.8 – 1.25. Overall, under plausible parametrizations the magnitude of the elasticity remains large.

Let the utility function be:

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

where ϵ is the intra-temporal elasticity of substitution between non-durable consumption and housing services.

Proposition 2 *Consider the model with CES preferences, liquid housing wealth, and $\theta = 0$. The individual response of non-durable consumption to an unexpected, permanent, proportional change in house prices dP/P is*

$$(\epsilon - 1) \frac{\frac{r+\delta}{1+r} PH_{it}}{C_{it} + \frac{r+\delta}{1+r} PH_{it}} C_{it} + MPC_{it} \cdot (1 - \delta) PH_{it-1}.$$

Proof. With constant prices, the user cost of housing (or implicit rental rate) is $\frac{r+\delta}{1+r}P$. So we can define total spending on non-durables and housing services

$$X_{it} \equiv C_{it} + \frac{r + \delta}{1 + r} PH_{it},$$

and the price index

$$P^X \equiv \left[\alpha^\epsilon + (1 - \alpha)^\epsilon \left(\frac{r + \delta}{1 + r} P \right)^{1-\epsilon} \right]^{\frac{1}{1-\epsilon}}.$$

The household's optimization problem can then be decomposed into an intertemporal optimization problem, characterized by the Bellman equation

$$V(W, s) = \max \frac{1}{1 - \sigma} \left(\frac{X}{P^X} \right)^{1-\sigma} + \beta E [V(W', s')],$$

subject to

$$W' = (1 + r) [W + Y(s) - X] \geq 0,$$

and an intratemporal utility maximization problem. The solution to the intertemporal problem is independent of P^X as it only appears as a multiplicative constant in the objective function. So the policy $X(W, s)$ is independent of P . The solution to the intratemporal problem gives

$$C = \frac{\alpha^\epsilon}{\alpha^\epsilon + (1 - \alpha)^\epsilon \left(\frac{r + \delta}{1 + r} P \right)^{1-\epsilon}} X, \quad H = \frac{(1 - \alpha)^\epsilon \left(\frac{r + \delta}{1 + r} P \right)^{-\epsilon}}{\alpha^\epsilon + (1 - \alpha)^\epsilon \left(\frac{r + \delta}{1 + r} P \right)^{1-\epsilon}} X.$$

The response of C to P conditional on X is then

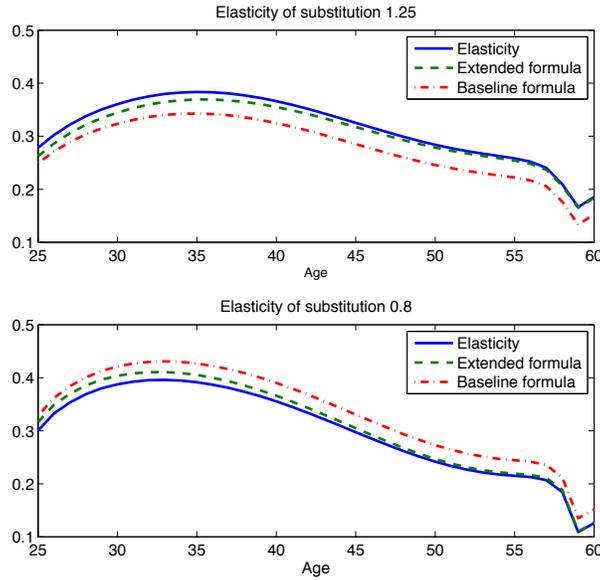
$$\frac{\partial C}{\partial P} = (\epsilon - 1) \frac{r + \delta}{1 + r} \frac{H}{X} C.$$

Combining this effect with the effect on X through W , yields the desired result. ■

An elasticity of substitution different from one implies that there is an additional term, proportional to the implicit share of housing services in the total consumption basket. In the model, this share is tightly linked to the ratio of house values to consumption. For example, take an agent with housing-to-consumption ratio 3.5—which is the roughly the average for agents in the 40s bin. With $r = 2.4\%$ and $\delta = 2.2\%$ this implies a share of housing services to total spending equal to 0.13. For such an agent if $\epsilon = 1.1$, the additional term is equal to $0.1 \times 0.13 = 0.013$. If the same agent has an MPC of 0.1 the magnitude of the baseline sufficient statistic is $0.1 \times 3.5 = 0.35$, so the additional term plays a minor role quantitatively.

For the case $\theta > 0$ analytical results are not available and we turn to simulations. Namely, for different values of ϵ we re-calibrate the model to hit the same targets as in the baseline and compute the elasticity of consumption to house prices. In Figure 19 we plot the computed elasticity for different age bins, together with our baseline formula and our amended formula. We present results for two cases which bound typical empirical estimates, $\epsilon = 0.8$ and $\epsilon = 1.2$. The amended formula provides a good approximation and the baseline still gives magnitudes close to the true responses.

Figure 19: CES



3 Extension: Stochastic House Prices

In this section, we explore the robustness of our baseline results (see Section 2.5) when house prices are stochastic. In particular, we assume that house prices follow a random walk with normal innovations. We calibrate this house price series using aggregate data from Corelogic from the period 1979-present. In this series the standard deviation of house prices is 0.05, the value we use. Other than allowing house prices to follow a random walk, we calibrate the stochastic price model to match the same targets as in the constant price model. See Section 2.4 for full details. The calibrated parameter values are shown in Table 2.

Table 2: *Parameter Values*

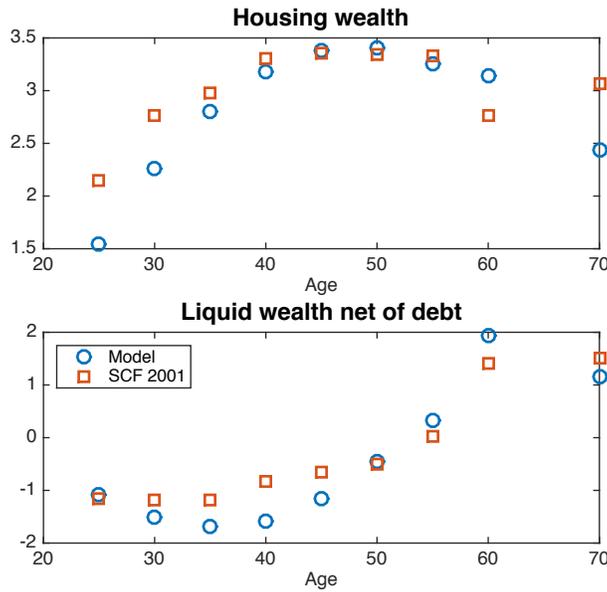
Parameter	σ	r	δ	θ	ρ_z	σ_z	α	β	Ψ	Ξ
Value	2	2.4%	2.2%	0.25	0.91	0.21	0.8521	0.9299	1,556	3.1002

Figure 20 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.⁶⁶ The blue circles are the model predictions and the orange squares are from the 2001 SCF. As in the baseline model, the stochastic price model delivers a reasonable fit to the data.

We now turn to the model predictions for consumption responses and look at the instantaneous response of non-durable consumption to a 1% shock to house prices.

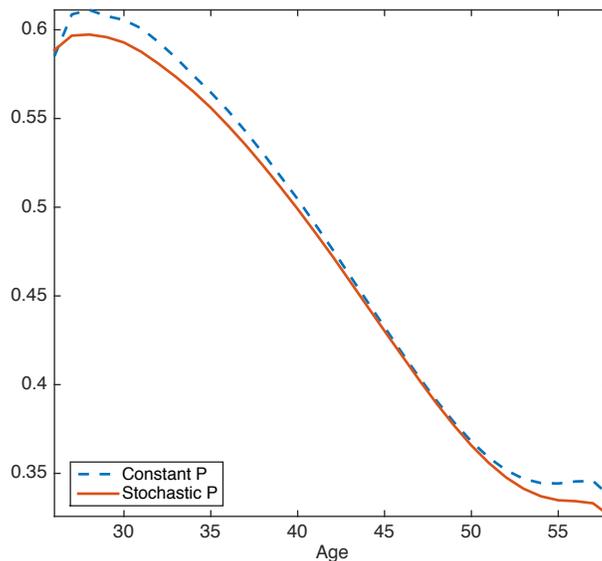
⁶⁶The point for the age bin 65 and over is plotted at age 70.

Figure 20: Housing and Net Liquid Wealth over the Life Cycle: Data vs Model



The results are shown in Figure 21. For ease of comparison, both the results from the baseline model (dashed blue line) and stochastic price model (solid orange line) are shown. The figure shows that the addition of house price risk barely changes the magnitude of the consumption elasticities. The average elasticity changes from 0.474 in our baseline model to 0.467 in the stochastic price model. Thus, allowing house prices to follow a random walk does not materially affect our main quantitative result.

Figure 21: Elasticities over the Life Cycle: Frictionless Model



4 Description of Computational Procedures

In this appendix, we describe the solution to the model with transaction costs and rental described in the body of the text. The household state vector is $\mathbf{s} \equiv (A, H, z, j)$, and the model is solved by backward induction from the final period of life. When working, households solve:

$$V(\mathbf{s}) = \max \{V^{adjust}(\mathbf{s}), V^{noadjust}(\mathbf{s}), V^{rent}(\mathbf{s})\}.$$

The three value functions, for adjusters, non-adjusters, and renters, are given by

$$\begin{aligned} V^{adjust}(\mathbf{s}) &= \max_{C, A', H'} U(C, H') + \beta E[V(\mathbf{s}')|z] \\ \text{s.t. } A' + PH' + C &= (1+r)A + Y(z) + (1-F)(1-\delta)PH, \\ A' &\geq -(1-\theta) \frac{1-\delta}{1+r} PH', \quad \mathbf{s}' = (A', H', z', j+1), \end{aligned}$$

$$\begin{aligned} V^{noadjust}(\mathbf{s}) &= \max_{C, A'} U(C, H) + \beta E[V(\mathbf{s}')|z] \\ \text{s.t. } A' + C &= (1+r)A + Y(z) - \delta PH, \\ A' &\geq -(1-\theta) \frac{1-\delta}{1+r} PH, \quad \mathbf{s}' = (A', H, z', j+1), \end{aligned}$$

$$\begin{aligned} V^{rent}(\mathbf{s}) &= \max_{C, A', \tilde{H}} U(C, \tilde{H}) + \beta E[V(\mathbf{s}')|z] \\ \text{s.t. } A' + C + \phi P\tilde{H} &= (1+r)A + Y(z) + (1-F)(1-\delta)PH, \\ A' &\geq 0, \quad \mathbf{s}' = (A', 0, z', j+1). \end{aligned}$$

The problem for a retired household is identical except that social security benefits replace labor earnings. At the age of retirement households also receive an additional lump sum transfer to match the level of retirement wealth which is now liquid, as described in the text. At the time of death households' continuation value is given by the bequest motive in the text.

To solve the model numerically, we proceed as follows. First, in order to rectangularize the choice set and simplify the computational problems imposed by the endogenous liquidity constraint, we follow Díaz and Luengo-Prado (2010) and reformulate our problem in terms of voluntary equity, defined as

$$Q \equiv A + (1-\theta) \frac{1-\delta}{1+r} PH.$$

After substituting the budget constraint into the utility function to eliminate C as a choice variable, the value function can then be rewritten in terms of the two non-negative state variables

Q' and H' . Note that A' and H' are chosen prior to next period shocks to house prices. Thus, shocks to house prices imply that realized voluntary equity will differ from the value chosen in the previous period. Namely, given a chosen pair Q', H' , the realized value of Q next period will be $Q' + (1 - \theta) \frac{1-\delta}{1+r} H' \Delta P$, where ΔP is the house price shock. This implies that although households are constrained to always choose $Q' \geq 0$, realized voluntary equity can be negative, for a large enough negative house price shock. To account for this, we solve the model for states that include negative voluntary equity even though households are constrained to choose non-negative values for this variable.⁶⁷

We discretize the problem so it can be solved on the computer by first discretizing z using the algorithm of Tauchen (1986). We use 13 grid points for z . We then approximate the functions V_j^{adjust} , $V_j^{noadjust}$, and V_j^{rent} as multilinear functions in the endogenous states. In our benchmark calculation, we use 120 knot points for Q (we space these points more closely together near the constraint) and 40 knot points for H . The presence of fixed adjustment costs on housing together with the borrowing constraint make the household policy function highly non-linear. For this reason, we follow Berger and Vavra (2015) and compute optimal policies for a given state-vector using a Nelder-Meade algorithm initialized from 3 different starting values, to reduce the problem of finding local maxima. The value of adjusting, not adjusting and renting are then compared to generate the overall policy function. We proceed via backward induction from the final period of life.

To simulate the model, we initialize cohorts to match the values of the SCF for age 22-27 year old households. First, we randomly split the sample into two groups to match the fraction of homeowners and renters. Then within each group, we split the sample into 4 income bins and assign the median value of housing and liquid assets from the SCF in that same income bin. (By definition, the value of housing for the renter groups is always zero). The model is simulated with 100,000 households and house price impulse responses are computed for each cohort.

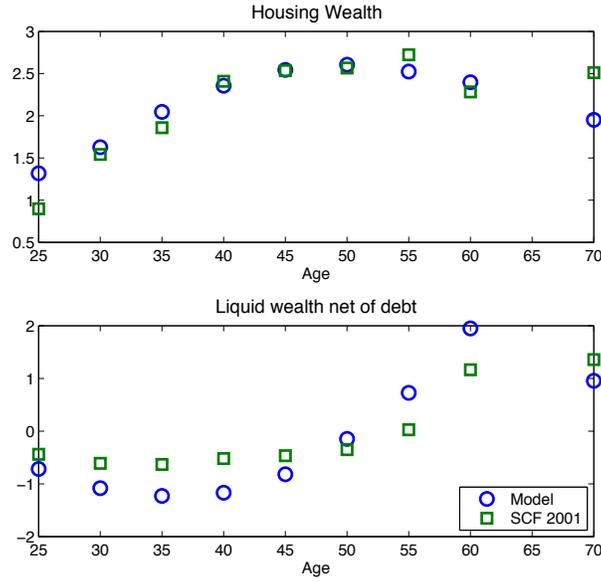
In the boom-bust episode, house price expectations become an additional state-variable, and the model is solved and simulated for the combination of each age-calendar time pair.

In section 3.3 we introduce two extensions to household mortgage contracts. We first introduce costly refinancing. In this version of the model, non-adjusting households can increase voluntary equity with no cost, as in the baseline model, but households who want to decrease voluntary equity must pay a fixed numeraire cost to do so. More specifically, $V(\mathbf{s}) = \max \{V^{adjust}(\mathbf{s}), V^{noadjust}(\mathbf{s}), V^{refi}(\mathbf{s}), V^{rent}(\mathbf{s})\}$ where $V^{noadjust}(\mathbf{s})$ now includes an additional constraint that $A' \geq A$, and $V^{refi}(\mathbf{s})$ is identical to $V^{noadjust}(\mathbf{s})$ in the baseline problem but with budget constraint $A' + C = (1 + r)A + Y(z) - \delta PH - F^{refi}$, where F^{refi} is calibrated at 0.005, to match targets described in the text.

In the second extension, we introduce long-term mortgages. Since interest rates in the model

⁶⁷Shocks to house prices in the model are not large enough to ever reach a situation where realized Q is so negative that households would be unable to choose $Q' \geq 0$ without having negative consumption.

Figure 22: Housing and Net Liquid Wealth over the Life Cycle: Data vs Model



are constant, long-term mortgages can be captured by changing the collateral constraint in $V^{noadjust}(\mathbf{s})$ to.

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

5 Additional Quantitative Results and Calibration Details

This section shows a number of numerical results and details which are suppressed from the main text for brevity.

Table 3 shows the best fit parameters for the model with transaction costs but no rental market, and Figure 22 shows the corresponding model fit.

Table 3: *Parameter Values for Model with Fixed Costs and No Rent*

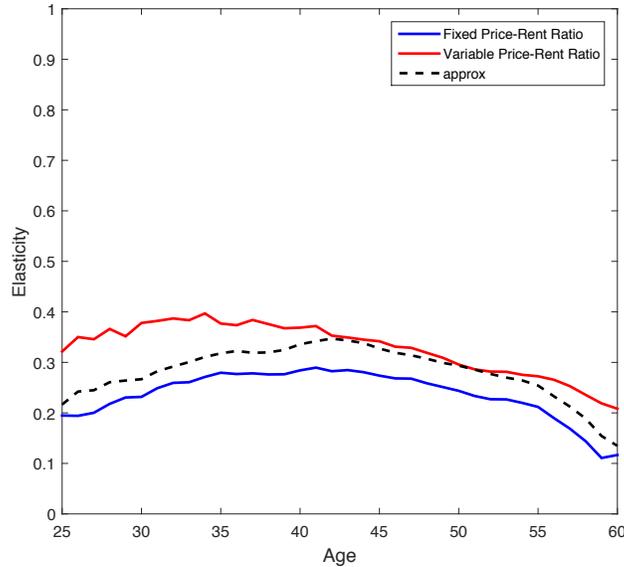
Parameter	σ	r	δ	θ	ρ_z	σ_z	α	β	Ψ	Ξ
Value	2	2.4%	2.2%	0.25	0.91	0.21	0.875	0.922	2,300	1.75

Table 4 shows the best fit parameters for the model with transaction costs and rental markets. See Figure 7 in the main text for corresponding model fit.

Table 4: *Parameter Values for Model with Fixed Costs and Rent*

Parameter	σ	r	δ	θ	ρ_z	σ_z	α	β	Ψ	Ξ	ϕ	$\hat{\phi}$
Value	2	2.4%	2.2%	0.25	0.91	0.21	0.86	0.922	2,300	2.25	.0605	.0472

Figure 23: Comparing Assumptions on Price-Rent Ratio

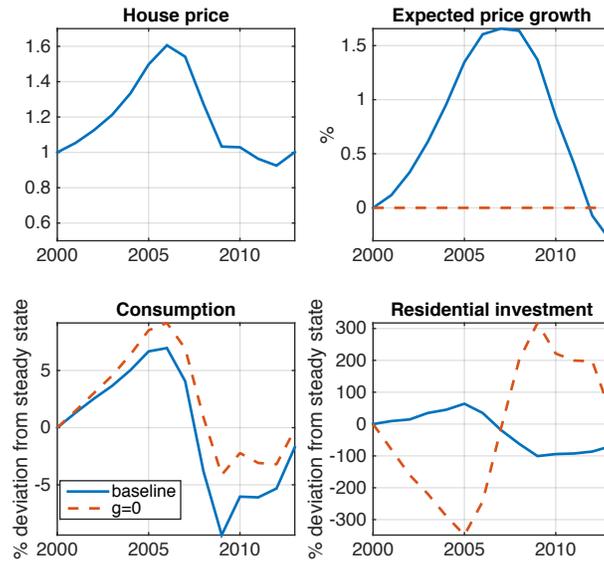


In most of our model results, we assume that the price-rent ratio is fixed so that rents change proportionately with house prices. It is straightforward to instead make the assumption that rents are fixed when house prices move so that the price-rent ratio moves one-for-one with house prices. Figure 23 compares the elasticities in the model with rentals under the constant and the variable price-rent ratio. It shows that elasticities are modestly amplified when the price-rent ratio moves with house prices. This is because if house prices go up but rents remain constant, homeowners can sell their houses, switch to renting and increase their consumption. It is also worth noting that the sufficient-statistic approximation continues to work well, although now it underestimates rather than overestimating the true elasticity.⁶⁸ In reality, the elasticity of rents to prices depends on the time-horizon but are neither completely fixed nor exhibit complete pass-through.

Finally, Figure 24 shows the boom-bust patterns for the frictionless model. Consumption responses are larger in the frictionless model, which is consistent with the larger elasticities derived in our partial equilibrium analysis. Another noticeable difference is that, with frictionless adjustment, residential investment is more sensitive to growth rate shocks, so fitting residential

⁶⁸Note that the approximation formula does not depend on how the price-rent ratio moves, in a model where households assume that prices are constant, neither MPCs nor housing values are affected by unanticipated changes in these ratios. Thus, there is only a single approximation line since it is equal in the two models.

Figure 24: Boom-Bust Simulation in Frictionless Model



investment movements in the data requires smaller g_t shocks (the g_t shocks in the full model with frictions and rentals look more in line with survey evidence).