Anticipated House Purchases and Household Consumption

Isaac Gross[1]

Monash University

October 20[th], 2018

Abstract

This paper evaluates how anticipated purchases of housing affect households' consumption of non-durable goods. I build a heterogeneous agent model in which households save in liquid assets and illiquid housing, where the latter can be used as collateral for borrowing. I show that it is able to replicate the empirical distributions of income and wealth within the US economy. In the model there is substantial heterogeneity in households' marginal propensities to consume. Households with a high probability of buying housing stock lower their consumption of non-durable goods in anticipation of being credit constrained after their purchase. This results in households with low, even negative, marginal propensities to consume. I verify the model's predictions using micro-data from the PSID to show that (i) consumption falls in anticipation of, and after, changes in the stock of housing and (ii) households who are planning on purchasing housing have negative marginal propensities to consume. Finally, I use this model to examine the general equilibrium effects of tax credits for first home buyers and show that they lead to decreases in aggregate consumption.

JEL codes: D14 , E21, E62.

Keywords: household consumption, heterogeneous households, housing demand, leverage.

1 Introduction

The recent financial crisis has highlighted the impact that household debt has on consumption. A burgeoning literature has documented how higher levels of household leverage were associated with deeper falls in consumption. Debt overhang has held back consumer spending and the subsequent recovery.[2]

In contrast to the unexpected deterioration of balance sheets during the crisis, most variation in household debt revolves around the decision to purchase a house, which is generally anticipated years in advance. In this paper, I show that these anticipated increases in debt lead to a decrease consumption of non-durable goods. I further show that policies aimed at increasing home ownership may therefore reduce overall consumption.

In this paper I make three contributions. First, I extend the multiple asset Aiyagari-Bewley-Huggett model to include a market for housing. My model builds on the two-asset framework from Kaplan, Moll and Violante (2016), by allowing households to use illiquid assets as collateral for borrowing. I show that households, anticipating that they will be credit constrained following the leveraged purchase of housing, reduce their consumption of non-durable goods in advance of, and after, their purchase.[3] This anticipatory saving causes households who have a high likelihood of purchasing housing to have negative marginal propensities to consume.[4] Second, I validate these results using micro-data from the Panel Study of Income Dynamics (PSID) to show that households who expect to adjust their ownership of housing in the next two years have lower, even negative, marginal propensities to consume. Finally, I use my model to estimate the impact of tax credits for first home buyers on non-durable consumption. I show that the tax credits cause aggregate consumption to fall as households have a larger incentive to save the necessary house deposit. After the tax credit expires, the share of highly leveraged households is elevated which futher surpresses consumption.

---

[2] Mian, Rao and Sufi (2013) and Mian and Sufi (2014) highlight how heterogeneity in households' balance sheets affected consumption dynamics during the recent recession. Carroll and Dunn (1997), Campbell and Cocco (2007), and Attanasio et al. (2012) also examine the relationship between the two using rich general equilibrium models. Finally, the empirical link between marginal propensities to consume and household leverage is explored by Broda and Parker (2014), Fagereng et al. (2016) and Kaplan et al. (2014).

[3] By contrast, in the standard two-asset model without leverage, the decision to buy illiquid assets triggers an increase in non-durable consumption (Kaplan and Violante, 2014).

[4] This anticipatory saving effect is distinct from precautionary saving, which occurs in response to uncertainty with respect to future income. The anticipatory saving effect exists in the absence of either aggregate or idiosyncratic uncertainty.

Structural models in which heterogenous agents save in a single risk-free asset are unable to match the empirical results on households' marginal propensity to consume .[5] Kaplan and Violante (2014) resolve this by modelling an economy in which households choose between a liquid and an illiquid asset, which is subject to a transaction cost. In their model, the existence of wealthy households that have few liquid assets generates higher marginal propensities to consume which are consistent with the empirical results. However their approach only models households' net asset position, and does not allow for borrowing. This assumption is incongrous with Misra and Surico (2014) who find that households with high levels of mortgage debt have a large propensity to spend.

My first contribution is to relax this assumption by developing a heterogeneous agent model à la Aiyagari-Bewley-Huggett, in which households save in either liquid assets or illiquid housing, and where housing can be used as collateral for liquid debt. Households save in order to self-insure against fluctuations in labour income, expand their access to credit and enjoy services from housing. Building on Caballero and Farhi (2014), Kaplan and Violante (2014) and Achdou et al. (2017), I introduce a housing market subject to several frictions. Households face borrowing constraints, financial frictions and transaction costs when trading houses. These frictions in the housing market are what generate a negative marginal propensity to consume for households that plan on purchasing a house. In a model without transaction costs, households optimise their consumption of housing and non-durable goods by keeping the marginal rate of substitution between the two constant. This implies a positive co-variance between consumption and housing, as income fluctuates over the life of the household. However, transaction costs induce households to make lumpy, debt-financed purchases of housing. When a household chooses to borrow in order to finance the purchase of housing stock, the opportunity cost of non-durable consumption in the present goes up for two reasons. First, the household faces a higher interest rate, since the rate they pay to borrow is higher then the rate they earned on their liquid savings. This increases the cost of consumption today relative to saving for consumption tomorrow. Second, borrowing to pay for housing increases the probability that households will be credit constrained in the future, which could force them to forego high-marginal-utility consumption. Therefore consumption today involves the possibility of foregoing a large amount of consumption utility tomorrow in the event that they are constrained.

---

[5]The collective evidence suggests a mean marginal propensity to consume of around 0.25.

The combination of a household having to pay a higher interest rate on their debt, and the desire to reduce the probability that they will hit their borrowing constraint, leads them to lower their level of consumption upon purchasing housing. Forward-looking households smooth this transition by lowering their level of consumption as the probability that they will choose to buy a house rises. In my model, a transitory rise in income increases the probability that a household will find it optimal to purchase of housing. This increases their incentive to save in anticipation of being credit constrained post-purchase. For the majority of households, the wealth and liquidity effects outweigh the anticipatory saving effect such that the increase in income leads to a rise in consumption. However, for households that are close to the point at which they would choose to purchase housing, a small increase in income can lead to a large rise in the probability that they will do so. For these households, the anticipatory saving outweighs the wealth and liquidity effects and thus produces a negative co-variance between income and consumption.

My second contribution is to estimate the effects of expected changes in household debt on non-durable consumption. I pursue two empirical strategies. First, I estimate households' marginal propensities to consume from transitory income shocks using the semi-parametric method developed by Blundell, Pistaferri and Preston (2008) and Kaplan and Violante (2010). I show that average marginal propensities to consume are lower for households that are more likely to buy additional stock, as measured by households' self-reported expectations that they will move in the next two years.[6] I find that households which report that they will definitely move have a marginal propensity to consume that is negative. Second, I estimate a series of models of consumption growth following the approach of Dynan (2012) to show that anticipated changes in household debt are associated with declines in household consumption in magnitudes that are consistent with the results from the model. I estimate these regression models using an instrumental variable approach which I show provides a consistent estimator of the impact of expected changes in debt and home ownership.

Finally, the existence of households with a negative marginal propensity to consume has important implications for public policy. During the financial crisis several governments implemented tax credits aimed at encouraging first home buyers to enter the market. My third contribution is to use my model to analyse the impact of these temporary tax credits on house prices, home ownership and aggregate consumption. I show that while the tax credits do boost prices and

---

[6]These expectations are surveyed by the PSID and are positively correlated with future changes in both home ownership and household debt.

turnover in the housing market, the effect is concentrated at the time the tax credit expires. The tax credit "pulls forward" sales that would otherwise have occurred in the future. By encouraging households to purchase more housing, it also increases the share of highly leveraged households within the economy which leads to lower aggregate consumption. This result contrasts with both the stated intention and the previous literature on such tax credits, which assumed that they would stimulate the economy and increase aggregate consumption.

The paper is organised as follows. Section 2 discusses how this paper contributes to the literature. Section 3 presents my general equilibrium model of the housing market under incomplete markets. Section 4 discusses the calibration of the model, analyses the implications for household consumption dynamics and discusses alternative calibrations. Section 5 validates the main findings of the model using micro-data from the PSID. Section 6 applies the model to analyse the impact of tax credits for first home buyers. Section 7 concludes.

2   Related Literature

My analysis relates to several strands of the literature.

First, I connect to the large theoretical literature that models how heterogeneous agents behave in the presence of incomplete markets. Introduced by Bewley (1983) and with early explorations by Huggett (1993) and Aiyagari (1994), these models explore how idiosyncratic shocks affect the distribution of wealth and income within the economy and how households achieve partial insurance by saving with risk-free assets. This approach solves for market-clearing prices by aggregating asset demand over the distribution of households. Further developments of this framework include extensions to multiple assets with different degrees of liquidity (Kaplan and Violante, 2014), and heterogeneity among household preferences (Iacoviello and Pavan, 2013). Previous literature on the two-asset model assumes that households are able to borrow up to an exogenous limit. I add to this literature by extending the two-asset model to allow for borrowing limits to be endogenously determined by a household's stock of housing and the market-clearing price.

More recent work by Achdou et al. (2017) shows that by solving the model in continuous time, aggregate shocks can be solved keeping the entire wealth distribution as a state variable. Guerrieri and Iacoviello (2017) highlight the importance of this approach, showing that aggregate shocks have quantitatively different impacts on agents in different regions of the distribution.

4

Second, my approach to modeling the mortgage market is based on studies focused on the housing markets. Given the computational complexity, life-cycle models typically assume that house prices are purely exogenous (Fernandez-Villaverde and Krueger (2011); Iacoviello and Pavan (2013); Yang (2009)). A handful of papers develop models with aggregate shocks that impact the equilibrium price of housing, but they require either environments with household heterogeneity limited to two agent models (Justiniano et al. (2015); Iacoviello and Pavan (2013)) or approximations which reduce their dimensionality to a single moment as pioneered by Krusell and Smith (1998). However, for more complex models the number of moments required for an accurate approximation can be orders of magnitude higher (Ahn et al., 2017). I add to this literature by solving a model which contains both an endogenously determined price of housing and a rich level of household heterogeneity capable of matching the empirical distribution of earnings.

This paper also contributes to the empirical literature studying the link between a household's balance sheet and non-durable consumption. A number of papers have used increasingly rich models with housing and debt to address aggregate questions and make cross sectional predictions, e.g., Carroll and Dunn (1997), Campbell and Cocco (2007), and Attanasio et al. (2012). The impact of exogenous house price movements on consumption, and debt over the life-cycle, has also been extensively examined using the heterogeneous agent framework, mostly notably by Garriga and Hedlund (2016), Berger et al. (2015) and Iacoviello and Pavan (2013). By contrast, my paper focuses on how transaction costs affect non-durable consumption before and after a household adjusts its stock of housing. Benmelech et al. (2017) and Best and Kleven (2013) both estimate the the impact of a house purchase on consumption. Their approach compares how consumption, as measured by the Consumer Expenditure Survey and the UK Living Costs and Food Survey respectively, varies with housing tenure.[7] They find that consumption of durable goods rises for households who have recently purchased a house, while consumption of non-durable goods falls. By contrast I use data from the PSID which follows households as they move and surveys households' expectations, at the cost of less comprehensive consumption data. I extend their analysis by using lagged household expectations as a instrumental variable to control for potential endogeneity within the model. This allows me to control for contemporaneous shocks which may affect both consumption and the decision to purchase a house. Another approach in the empirical

---

[7]Since neither survey follows households that move, Benmelech et al. (2017) and Best and Kleven (2013) focus how consumption is affected after a house purchase.

literature is Martin (2003), who shows that households increase their spending on food before moving to a smaller house and decrease it before moving to a larger house. I replicate this analysis with data on broader consumption and extend it to measuring a household's marginal propensity to consume just prior to moving.

This paper also contributes to the fast expanding literature that estimates heterogeneity in marginal propensities to consume across households. There is a wealth of literature that analyses how debt affects households' marginal propensities to consume, using self-reported values (Auclert, 2017), exogenous variation in fiscal transfers (Greer, Parker and Souleles (2006) and Parker et al. (2013)), lottery winnings (Fagereng et al., 2016) and a semi-structural approach (Blundell, Pistaferri and Preston 2008). These methods consistently find that households with high levels of household debt and low access to liquid assets have high marginal propensities to consume. I contribute to this literature by using data on households' self-reported expectation of moving to show that households with a high probability of buying additional housing have lower, even negative, marginal propensities to consume.

Finally, in an application of the model I estimate the impact of tax credits for first home buyers. Dynan et al. (2013) find that while there is some evidence that the tax credit helped support house prices, there was no discernible positive impact of the tax credits on broader economic activity. My work suggests these tax credits have a contradictory impact on aggregate consumption, potentially resolving this puzzle.

3   The Model

In this section I introduce a two-asset, incomplete markets model. My main innovation is to augment Kaplan, Moll and Violante (2016)'s rich representation of household consumption and saving behavior with the ability of households to borrow against the nominal value of their stock of illiquid housing. Households face uninsurable idiosyncratic shocks to labour income and can self-insure through saving in either housing and liquid assets. Outside of the household and housing sector, the rest of the model is kept as parsimonious as possible. To economise on computational time I solve the model in continuous time using the upwind approximation outlined in Achdou et al. (2017).

3.1 Households

The economy is populated by a continuum of households who are differentiated by their holdings of liquid assets b, their ownership of illiquid housing a, and their idiosyncratic labour productivity z. I assume that time is continuous and that there is no aggregate uncertainty, thus at each point in time t the economy is governed by the joint distribution of the households $\mu_t(b, a, z)$. Households die off with a fixed probability λ and are replaced by new households with zero net wealth $(a = b = 0)$ and the mean level of income.[8]

Households must optimise the expected present value of their utility flow, $u(c_t, h_t, l_t)$, by choosing the optimal path of their of labour supply, l_t, non-durable consumption c_t and their ownership of housing h_t. Each household takes prices, wages, interest rates and governmental transfers, $\Phi = \Phi\{w_t, r_t^b, a_t^p, \tau_t, T_t\}$, as given. Preferences are time-separable, and conditional on surviving, the future is discounted at a rate $\rho \geq 0$:

$$U_t = E_t \int_t^\infty e^{-(\rho+\lambda)t} u(c_t, h_t, l_t) \, dt. \tag{1}$$

The flow utility is increasing and strictly concave in c and h and decreasing and strictly convex in l. I assume the functional form

$$u_t = \frac{\left\{ \{c_t - z_{it}G(l_t)\}^{1-\zeta} \{h_t\}^\zeta \right\}^{1-\gamma} - 1}{1-\gamma} \quad G(l_t) = \phi z_t \frac{l_t^{1+\frac{1}{\eta}}}{1+\frac{1}{\eta}}, \tag{2}$$

where the curvature parameter γ determines households' risk aversion and inter-temporal elasticity of substitution, ζ is the weight placed on housing relative to non-durable consumption, and η is the Frisch elasticity of labour supply. I assume a functional form for flow utility that follows Greenwood et al. (1988), which ensures there is no wealth effect on labour supply, l_t. Thus all households, regardless of their assets, will supply the same amount of labour which is solely a function of the wage per effective unit of labour, w_t, and does not vary with idiosyncratic labour productivity shocks, z_t.[9]

Households maximise the present value of expected utility subject to four constraints. The first is the budget constraint, which governs the evolution of liquid assets, b_t, in the form of a

---

[8]The stochastic death of households is not required to solve the model, but is critical to matching the high share of renters with no liquid assets observed within the economy.

[9]This assumption creates a direct mapping between productivity and earnings, which facilitates the calibration of the exogenous productivity process. I relax this assumption in Section 5.

risk-free bond

$$\dot{b}_t = z_t w_t l_t + r_t b_t - c_t - c_t^{rent} p_t^{rent} - k\left(a_t, a_t'\right) - p_t^a \Delta a_t - T_t \tag{3}$$

Housing assets are illiquid as households must pay a fee $k\left(a_t, a_t'\right)$ in order to buy or sell an amount of housing. I set this friction as a fixed percentage of the current level of housing , a_t, and new level of housing, a_t'. I denote any new purchases or sales of housing $\Delta a_t = a_t' - a_t$. Note that while the model is written in continuous time, any purchase of housing stock and the accompanying change in liquid assets occur instantaneously. Each household's stock of housing depreciates over time at a rate δ, and thus the stock of housing evolves according to

$$\dot{a}_t = (1 - \delta) a_t + \Delta a_t \tag{4}$$

The interest rate paid on the liquid asset varies depending on whether the household is a saver, who deposits their funds with the financial firms, or a net borrower. The spread between saving and borrowing occurs due to the cost of financial intermediation and is detailed below.

$$r_t = \begin{cases} r_t^+ & \text{if } b_t > 0 \\ r_t^+ + r_t^{spread} & \text{if } b_t \leq 0 \end{cases} \tag{5}$$

Each household's consumption of housing services is defined by the stock of housing that they own. Households that do not own housing must purchase housing via the rental market. However, rental units only grant a fraction, θ^r, of the benefit to households relative to owned stock.

$$h_t = \begin{cases} \theta^r c_t^r & \text{if } c_t^r > 0 \\ a_t & \text{if } c_t^r \leq 0 \end{cases} \tag{6}$$

Finally, each household can borrow against a proportion of the total value of the stock of housing they own.

$$b_t \geq -M a_t p_t^a \tag{7}$$

$$a_t > 0 \tag{8}$$

Households maximise 1 subject to (3)-8, taking as given the path for the real wage $\{w_t\}_{t\geq 0}$, the interest rate on liquid assets $\{r_t^b\}_{t\geq 0}$, the price of housing relative to non-durable goods $\{p_t^a\}_{t\geq 0}$, and taxes and fiscal transfers $\{\tau_t, T_t\}_{t\geq 0}$. As described below, the paths for the price variables are determined by the market-clearing conditions for labour, liquid assets and housing. In Appendix A I describe how the households' optimization problem can be solved recursively with a Hamilton-Jacobi-Bellman equation. The resulting steady state solution yields decision rules governing household behavior as a function of liquid assets, illiquid housing, productivity and relative prices for non-durable consumption $c(a, b, z; \Phi)$, labour supply $l(a, b, z; \Phi)$, and the housing adjustment rule $\alpha(a, b, z; \Phi)$. Outside of the steady state these rules will be time-varying and respond to the time path of aggregate prices $\Phi_t = \{w_t, r_t^b, p_t^a, \tau_t, T_t\}_{t\geq 0}$. These optimal decision rules imply drift paths for households' holdings of liquid assets and housing, which when combined with the exogenous stochastic process for z, can be used to calculate the stationary joint distribution of assets and income $\mu_t(a, b, z, \Phi)$. In Appendix A I also outline the Kolmogorov forward equation that defines the evolution of this distribution over time.

3.2 Firms

Final Goods Firms

A continuum of final goods firms produce goods for consumption, $y_{j,t}$, by using effective units of labour $n_{j,t}$ in a linear production function. These firms are perfectly competitive, take the real effective wage, w_t, as given and have zero profit. Their prices are perfectly flexible and are defined as the numaire. They maximise profits

$$\Pi_t^{goods} = y_t - w_t n_t, \tag{9}$$

subject to the production function

$$y_t = \zeta_t n_t, \tag{10}$$

where ζ_t is an aggregate labour productivity shock.

Real Estate Firms

Real estate firms borrow from financial firms to fund the purchase of housing, which is rented out to households that do not own their own home. Real estate firms operate in a perfectly competitive market, and thus the rental price is pinned down by the cost of borrowing, a fixed cost of providing rental services, the price of purchasing house stock and any expected capital appreciation driven by the drift in the price of housing. I assume that real estate firms are able to borrow against the full value of their housing stock. The zero profit condition pins down the price of rental stock as the function of the cost of housing, capital gains from housing and the interest rate

$$p_t^{rent} = \phi^{rent} + \left(r_t^{b-} + \delta \right) p_t^a - \dot{p}_t^a, \tag{11}$$

where ϕ^{rent} is the fixed cost of providing rental services.

Financial Firms

Financial firms intermediate saving and borrowing within the economy subject to a fixed cost, r_t^{spread}, in the form of a spread between the interest rate on deposits and loans. Households make liquid deposits in financial firms, which are then loaned out to households who require a mortgage or to real estate firms. Finally, I assume that these financial firms operate in a competitive environment and thus earn zero profits

$$\Pi_t^{financial} = \left(r_t^b + r_t^{spread} \right) \left(\int_{b<0} 1 \, d\mu_t + \int_{a=0} c_t^{rent} \, d\mu_t \right) - r_t^b \int_{b>0} d\mu_t. \tag{12}$$

Construction Firms

Finally, a continuum of construction firms produce new units of housing for sale to both households and real estate firms. These firms are perfectly competitive and produce houses using labour, n_t^a, and housing permits, $\overline{L}$, purchased from the government. This ensures that the price elasticity of supply of housing is finite, and is akin to assuming investment adjustment costs in the housing sector. They maximise profits

$$\Pi_t^{construction} = p_t^a I A_t - w_t n_t^a - p_t^L \overline{L} \tag{13}$$

subject to the production function

$$IA_t = \zeta_t^a \, (n_t^a)^\omega \, \bar{L}^{1-\omega}. \tag{14}$$

Thus the price elasticity of supply is given by $\frac{\omega}{1-\omega}$.

I follow Favilukis et al. (2017) by assuming that the government sells the housing permits at the market-clearing price. This ensures that all rents from the fixed supply of housing permits accrue to the government and the construction sector makes no profit in equilibrium.

Maximizing 13 subject to 14 generates the housing investment function

$$IA_t = (\omega p_t^a)^{\frac{\omega}{1-\omega}} \, \overline{L} \tag{15}$$

3.3 Government

The fiscal authority funds an exogenous level of government spending and fiscal transfers to households via a combination of lump-sum payments and taxes on labour income.

$$T_t = \tau_t w_t l_t + \tilde{T}_t \tag{16}$$

Any fiscal deficit must be financed by issuing debt to households in the form of liquid assets. I assume that the government are able to borrow directly from savers and are not subject to the friction generated by the financial intermediaries. Thus the government's budget position is given (in deficit terms) by

$$\dot{B}_t^{gov} = r_t^b B_t^{gov} - \int T_t \, d\mu_t - p_t^L \bar{L} + G_t \tag{17}$$

3.4 Market Clearing Conditions

An equilibrium in this economy is defined as paths for individual household and firm choice variables $\left\{ a_t, b_t, c_t, c_t^{rent}, \alpha_t, l_t, n_t, n_t^a \right\}_{t \geq 0}$, and aggregate prices $\left\{ w_t, r_t^b, p_t^a, \tau_t, T_t \right\}_{t \geq 0}$ such that households and firms optimise their objective functions subject to their respective constraints, taking prices as given, for all time t and where the distribution of households, $\mu_t \, (a_t, b_t, l_t)$, is such that (i) the goods market, (ii) the market for liquid assets, (iii) the market for housing and (iv) the labour market all clear.

The final goods market clears if the total production by final goods firms is equal to consumption, government services, and the total resources lost to the frictions related to the housing sector, the financial firms and the real estate sector.

$$Y_t = \int_{a_t \neq a'_t} k\left(a_t, a'_t\right) d\mu_t + \int_{b<0} r_t^{spread} d\mu_t + \int_{a=0} \phi^{rent} d\mu_t + C_t + G_t \tag{18}$$

The liquid asset market clears when net household savings, B_t^{net}, are equal to the funds borrowed by the real estate sector and the fiscal authority.

$$\int b \, d\mu_t = B_t^{net} = p_t^A C_t^{rent} + B_t^{gov} \tag{19}$$

The market for housing must clear when total inflows from additional construction are equal to the aggregate depreciation plus the change in total demand by both home owners and renters.

$$IA_t = \delta A_t + \int \dot{a}_t + \dot{c}_t^{rent} d\mu_t \tag{20}$$

Finally, the labour market clears when the total labour demanded by the final goods and construction sectors equals the effective supply from households.

$$N_t + N_t^c = \int zl\left(a, b, z\right) d\mu_t \tag{21}$$

4   Calibration

Household Preferences

I calibrate the household utility function to match key moments in the data. I set the disutility of labour ϕ such that hours worked is equal to $1/3$ of the time endowment in equilibrium. I set the weight on housing in the utility function to match the average expenditure share on housing at 0.16. I set the curvature parameter γ, which governs the risk aversion and inter-temporal elasticity of substitution, to 1. Given these values, the average value for the inter-temporal elasticity of substitution is approximately 0.7.[10]

---

[10]The inter-temporal elasticity of substitution is defined as $\frac{-u_c}{c.u_{cc}}$. Given our assumption on the utility function, this corresponds to $\frac{c-G(l)}{(\gamma+\zeta(1-\gamma))c}$ which varies across the distribution of households.

Finally, I set the rate of household deaths λ to $\frac{1}{180}$ such that the expected lifespan for a newly created household is 45 years.

Income Process

A critical input into the distribution of liquid and illiquid assets is the frequency and size of earnings shocks that households are subject to. An environment in which incomes are subject to small, but frequent, shocks will encourage households to hold a high level of liquid stocks to better smooth consumption. By contrast when shocks are infrequent, but large, households will be more willing to hold illiquid assets, paying the adjustment cost only occasionally when a large shock occurs. It is thus important to match the higher moments of the earnings process faced by households if we are to replicate the empirical distributions in the data. One advantage of the continuous time environment is that we can calibrate shocks by their frequency of arrival, size and persistence. Conversely in a discrete time model shocks are assumed to arrive at the rate of one per unit of time.

As shown by Guvenen et al. (2015) the distribution of the change in log-earnings has a high degree of kurtosis (Table 1). In order to replicate this non-Gaussian distribution, I follow Kaplan et al. (2016) in assuming that household earnings follow a "jump-diffusion" process. This process allows us to generate a distribution for changes in log-earnings which match the high levels of kurtosis seen in the data.

Since all workers face the same wage per effective unit of labour and choose to supply the same amount of labour, a set of moments of households labour income can be used to estimate the labour productivity process z_{it}. I split the aggregate log-earnings process, z_{it}, into two components

$$z_{it} = z_{1,it} + z_{2,it}, \tag{22}$$

where each component, $z_{j,it}$, evolves according to

$$dz_{j,it} = -\beta_j z_{j,it} dt + dJ_{it}, \tag{23}$$

where dJ_{it} is the jump process which arrives at a rate α_j, such that over a small time period dt the probability that a jump occurs is $\alpha_j dt$ and the probability that a jump does not occur is

$(1 - \alpha_j) \, dt$. Conditional on a jump occurring, the new level for the component, $z_{j,it}$, is drawn from a normal distribution with zero mean and variance σ_j^2. Thus

$$dJ_{j,it} = -z_{j,it} + \epsilon_{j,it} \quad \text{with} \quad \epsilon_{j,it} \sim N\left(0, \sigma_j^2\right)$$

To calibrate the parameters of the earnings process I use a simulated method of moments to match data from Guvenen et al. (2015).

Table 1: Earnings Process Moments

Moment	Data	Process - Continuous Time	Process - Discreteised Grid
Variance - log earning	0.70	0.70	0.72
Variance - 1yr change	0.23	0.23	0.21
Variance - 5yr change	0.46	0.46	0.48
Kurtosis - 1yr change	17.8	16.5	17.3
Kurtosis - 5yr change	11.6	12.1	11.6
Share of 1yr change $< 10\%$	0.54	0.56	0.59
Share of 1yr change $< 20\%$	0.71	0.67	0.61
Share of 1yr change $< 50\%$	0.86	0.85	0.89

The estimates of the parameters suggest the presence of two kinds of shocks. The first is an infrequent, but persistent career shock which occurs on average every 38 years and has a half-life of 18 years. The second is a more temporary shock which arrives roughly every 3 years, but largely dissipates within a quarter.

Table 2: Earnings Process Parameters

Parameter	Persistent Component $j = 1$	Transitory Component $j = 2$
Arrival rate	0.08	0.007
Mean reversion	0.76	0.009
St deviation of jump process	1.76	1.54

Note. Rates are expressed as quarterly values.

4.1 Housing Market

I calibrate the parameters related to the housing market to match key long-run relationships in the US housing market. The fixed cost faced by households when adjusting their housing stock

is set to 7 per cent of the value of the house price. This produces a housing market in which 9.2 per cent of the housing stock is turned over annually. This compares with roughly 10 per cent in the data as estimated by Ngai et al. (2016).

The elasticity of housing supply is determined by the exponent in the production function of the construction sector, ω, which I set to generate an elasticity of 1.5 per cent, which is the median value of elasticities estimated by Saiz (2010) across US cities. I calibrate the size of the construction sector to be 5 per cent on total output within the economy, matching its long-run share of gross value added.

The remaining parameters $\left(p^a, \theta^r, \rho, s^h, \phi^{rent}\right)$ are set to match five key moments in the data.

Table 3: Calibrated Moments

	Model	Data
Share of owners with a mortgage	0.55	0.57
Share of renters	0.38	0.38
Mean net worth - income ratio	5.2	5.5
Average ratio earnings owners - renters	2.1	2.2
Median loan-to-value ratio	0.66	0.63

We match these moments by simulating the model over a grid of parameters, and choosing the set which minimises the difference between the five parameters.

Given that the majority of lending within the model occurs between savers and mortgagors, I calibrate the financial friction to match the long-run average spread between 30-year fixed rate mortgages and the 3-month Treasury Bill. Since 1985 this spread has averaged 3.5 per cent. Finally, the real interest rate and the real wage are internally calibrated to clear the labour markets.

4.2 Equilibrium

The households' optimal consumption and saving decisions and the evolution of the joint distribution of their income, housing and liquid assets can be summarised by a Hamilton-Jacobi-Bellman equation and a Kolmogorov Forward equation. The method of approximating these equations is outlined in Appendix A.

How well does this model match the distribution of wealth within the economy? In Table 4 I show some key moments related to the ownership of housing and total household net worth in the model.

Table 4: Selected un-targeted moments of the wealth distribution

	Data	Model
Housing owned by top 1%	19.43%	8.5%
Housing owned by top 10%	64.7%	55.3%
Housing owned by top 20%	79.4%	73.0%
Gini coefficient - housing	0.73	0.65
Net worth owned by top 1%	35%	19.5%
Net worth owned by top 10%	88%	76.3%
Net worth owned by top 20%	96%	91.2%
Gini coefficient - total net worth	0.85	0.79
Aggregate debt to house value	0.42	0.40

Source: 2004 Survey of Consumer Finance

On the whole, the model matches the data fairly well with the notable exception of the top end of the distribution. While the model generates significant degrees of inequality across both liquid assets and housing, it fails to match the large shares of wealth owned by the top 1 per cent of households.

4.3 Marginal Propensities to Consume and Household Leverage

How well does this heterogeneous approach to households match the empirical findings on households' marginal propensities to consume? Of the range of empirical studies that examine this topic, the most compelling use exogenous variation in fiscal payments (Kaplan and Violante (2014); Broda and Parker (2014)), and lottery winnings (Fagereng et al. (2016)). These studies suggest that households spend 15-25 per cent of one-off payments in the quarter that they are received.

A marginal propensity to consume can be thought of as the impact of a one-off increase in liquid wealth. I thus define the marginal propensity to consume of a payment of x additional dollars over a length τ periods as

$$MPC_\tau^x (a, b, z) = \frac{C_\tau (a, b + x, z) - C_\tau (a, b, z)}{x},$$

where $C_\tau (a, b, z)$ is the sum of expected consumption of an individual household over the next τ periods.

$$C_\tau (a, b, z) = E \left[\int_0^\tau c (a_t, b_t, z_t) \, dt | a_0 = a, b_0 = b, z_0 = z \right]$$

16

This formula can be solved using the Feynman-Kac formula as outlined in Appendix B.[11]

The unconditional quarterly marginal propensity to consume in the model is 19 per cent, which is consistent with the range of empirical results. However this unconditional mean masks a high level of heterogeneity across the distribution of wealth among households. Figure 1 shows the distribution of the marginal propensities to consume as a function of households' liquid and illiquid assets.

Figure 1: Household Heterogeneity
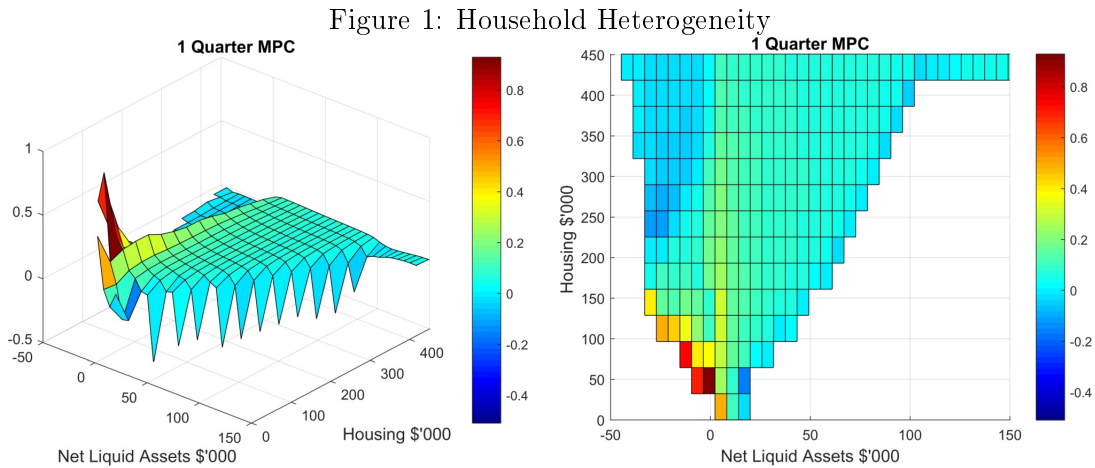
Note. This figure shows the distribution of marginal propensities to consume across households from two different viewpoints, over the subsequent quarter for households with the median income. For a 2-dimensional slice of this distribution see Figure 3

There are three notable features of the distribution in Figure 1. The first is the spike as households approach the borrowing constraint which is governed by the nominal value of their housing stock. The second is the high marginal propensity to consume for households that own some amount of housing stock but have zero liquid wealth. The third notable feature is the fall in marginal propensities to consume, often to negative levels, of households who are close to the point at which they would optimally choose to adjust their stock of housing.

The first two features are well documented in previous research. Broda and Parker (2014) find that households with limited access to liquid funds had a significant increase in consumption in response to the 2008 fiscal stimulus payments. In addition, Kaplan, Violante and Weidner (2014) use a semi-structural approach to measure marginal propensities to consume and find that households with limited access to liquid assets have significantly higher marginal propensities to

---

[11]An alternative strategy would be to estimate the impact of a small increase in the transitory component of the households' productivity process. This alternative strategy produces qualitatively similar results, but with larger simulation error due to the coarser grid used to approximate the productivity process.

consume, even if they have high levels of illiquid assets. These results are consistent with the existence of wealthy hand-to-mouth households. Fagereng et al. (2016) examine the response of household consumption to exogenous lottery prizes using Norwegian administrative data. They find that marginal propensities to consume vary with the amount of households' liquid assets and that households with close to zero liquid assets have high even if they are wealthy in terms of their illiquid asset position.

The existence of households with a negative marginal propensities to consume is a more novel finding. The key to understanding this result is examining how a household's decision to adjust its stock of housing affects its incentives to consume non-durable goods.

Since changes in a household's ownership of housing is subject to transaction costs, it will be optimal for households to make lumpy purchases, waiting longer before deciding to interact with the market and adjusting by a larger amount when they do so. These large, lumpy purchases require financing in the form of mortgages, which will place post-purchase households closer to the endogenous credit constraint and subject to the higher interest rate faced by borrowers relative to their previous position as savers.

These two effects are partially offset by the complementarity between housing and the consumption of non-durable goods in the households' utility function. However this effect is quantitatively small.

The combination of a higher probability of being credit constrained and the increase in interest rates faced by the household will increase the incentive for households to repair their balance sheet by reducing spending and paying back debt. Thus, on average, households decrease their consumption by 6 per cent conditional on buying additional housing stock (Figure 2).

Prior to purchasing additional housing stock, households anticipating the impending increase in debt and corresponding fall in non-durable consumption increase their savings rate to better smooth consumption in future periods. This anticipatory saving effect is stronger the higher the probability that households will find it optimal to buy a house in the near future.

Thus when households receive a small increase in liquid wealth there are three effects on household consumption. The wealth and liquidity effects both lead to an increase in the rate of consumption. However due to the increase in the probability that a household will choose to buy a house in future periods, the anticipatory saving effect will cause households to decrease their consumption. For the vast majority of households, the first two effects outweigh the third

Figure 2: Change in non-durable consumption conditional on buying additional housing stock
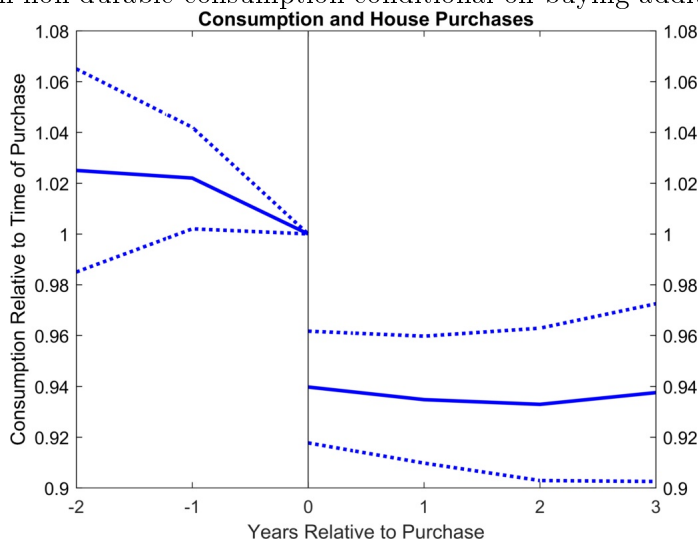
Note. The sold line represents the average level of consumption conditional on a household buying additional housing, relative to the level just prior to the purchase. The dashed lines represent the 90 per cent confidence interval.

as only a relatively small share of households adjust their housing stock annually.

The same logic holds for households that are close to selling a portion of their housing stock. When a household decides to reduce its stock of housing, it uses the proceeds from the sale to pay back its mortgage debt, thus reducing the probability that it will become credit constrained in the future. This increase in liquid assets will lead to the household increasing its non-durable consumption after reducing its stock of housing. Anticipating this increase in consumption, households will optimally choose to increase their consumption as their portfolio of assets approaches the adjustment point. Since a small increase in liquid assets moves such a household further away from the adjustment point, this will result in the household lowering their level of consumption as the probability that they will sell a portion of their housing stock decreases.

While this anticipatory saving effect reduces households' marginal propensity to consume, whether they adjust their stock of housing higher or lower, the size of the effect is asymmetric. This is because the transaction costs associated with adjusting the stock of housing reduces a household's wealth, regardless of whether they are buying or selling. This decrease in wealth ensures that the decrease in consumption associated with the purchase of additional housing stock is larger than the increase in consumption associated with selling the same amount. Thus

19

the anticipatory saving effect is stronger for households who expect to buy housing compared with those who expect to sell.

This effect is illustrated in Figure 3 which shows how a household's consumption varies with respect to liquid assets for a given level of housing and income. If the households stock of liquid assets becomes too high (low) it buys (sells) housing stock which results in a fall (rise) in non-durable consumption post-adjustment. The household anticipates the possibility of this adjustment occurring and begins to smooth consumption towards the post-adjustment level. This results in consumption falling (rising) as liquid wealth rises (falls) close to the optimal adjustment point which in turn implies that the marginal propensity to consume is negative.

Figure 3: Consumption as a function of liquid assets
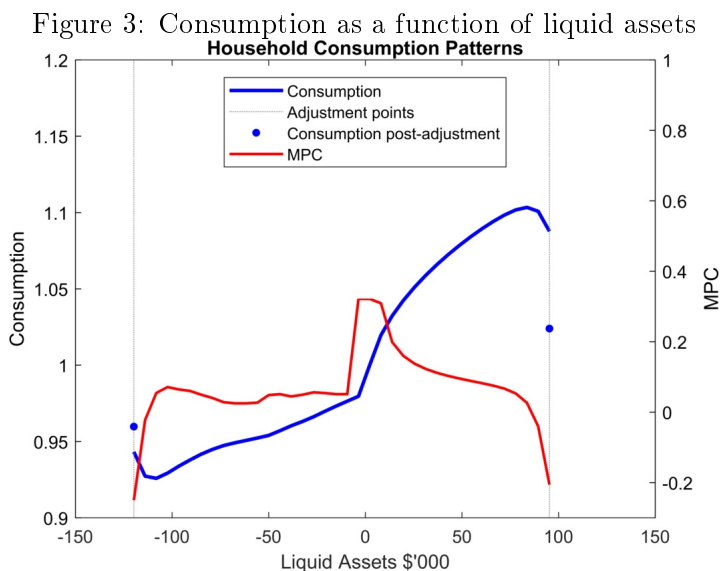
Note. This figure shows consumption as a function of the stock of liquid assets holding income and the stock of housing constant.

In practice, this phenomenon is caused by the combination of the cost of adjusting housing stock and the interest rate spread between borrowers and savers. When renters decide to enter the housing market they change from being savers to borrowers. This increases the interest rate they face when their liquid asset position shifts and causes them to decrease their level of non-durable consumption. If the stock of housing can be frictionlessly adjusted, then households are able to continuously adjust their ownership of housing and any debt as needed. A household's ability to slowly change the mix of their asset portfolio eliminates the large jumps in debt that are necessary in the presence of adjustment costs and allows consumption of both housing and non-durable goods to increase as household wealth and income increase. Alternatively, if the

financial friction is removed then the interest rate households face remains unchanged, even after they take out a mortgage reducing the change in consumption.

Why would households choose to lower their consumption as they anticipate adjusting their housing stock and deviate from the optimal consumption path imposed by the short-run Euler equation? Households could alternatively choose to buy additional housing when they have sufficient liquid wealth to avoid going into debt and risk becoming credit constrained. The answer is that households are better off financing their house purchases with debt because the alternative entails either (i) paying the transaction cost more often as households have less scope to expand their stock of housing; or (ii) remaining with an inefficiently small amount of housing as households save sufficient liquid wealth to be able buy more; or (iii) having to reside in inferior rental stock for longer which is subject to mark-ups by real estate firms.
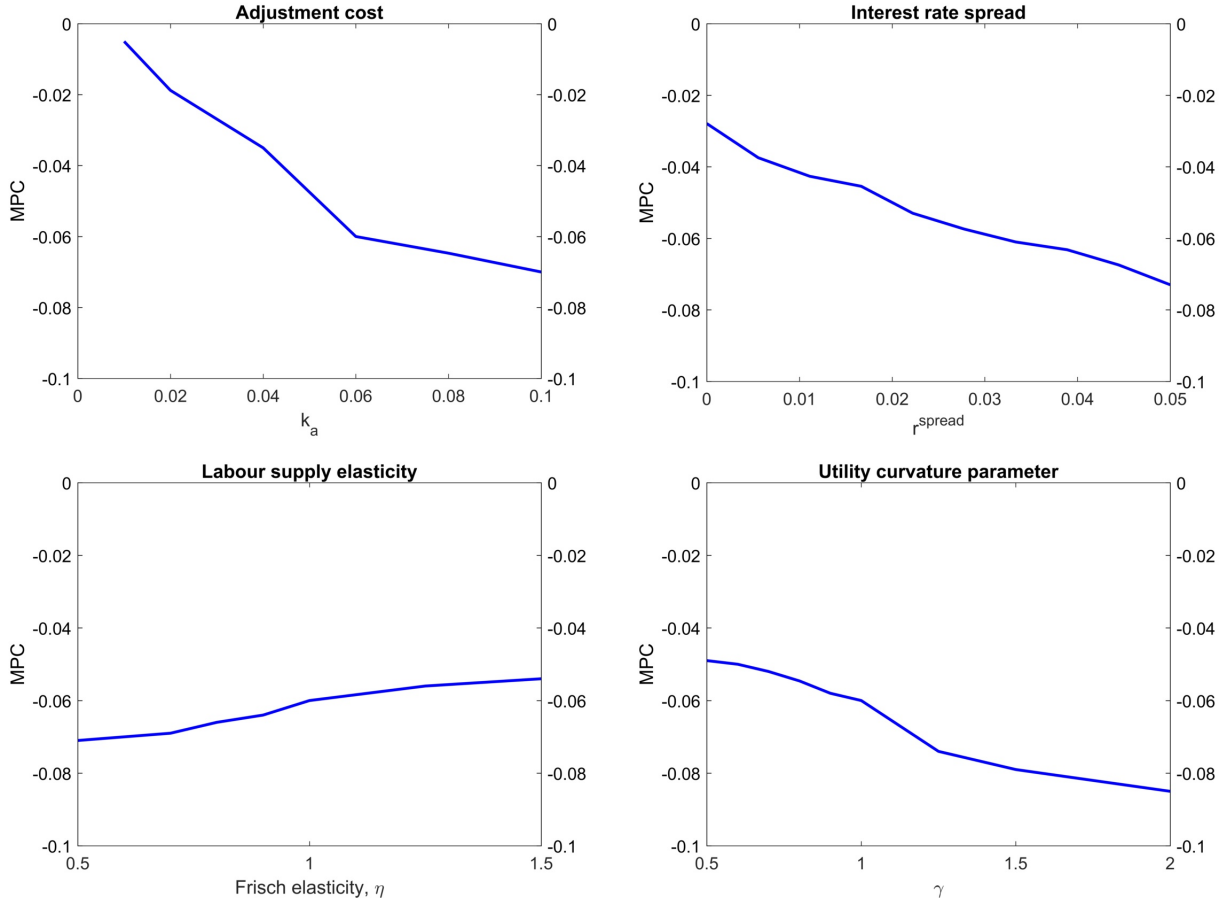
4.4 Robustness

To test the robustness of this result I estimate the model varying a range of externally calibrated parameters. Figure 4 shows how the main result, the negative marginal propensity for households close to their adjustment point, varies with the externally calibrated parameters. I measure this result by calculating the mean marginal propensity to consume for households that have a high likelihood of adjusting their stock of housing in the next year.

The existence of households with a negative marginal propensity to consume is relatively robust over a range of calibrations. However, the importance of adjustment costs and financial frictions are highlighted by these results. As these parameters are lowered, the proportion of households with negative marginal propensities to consume tends towards zero. As the adjustment cost decreases households are able adjust their stock of housing more frequently, as their wealth increases, without going into debt. The reduction of leveraged purchases decreases the proportion of households that will decrease their consumption when they buy housing stock.

In Appendix C I explore two different model specifications the first allows for the debt constraint to apply only when households adjust their stock of housing. At all other times households are able to keep the stock of debt constant, even if a fall in house prices lowers the value of their housing stock. This breaks the tight link between house prices and household debt as discussed by Iacoviello and Pavan (2013). The second model relaxes the assumption that household labour supply is identically supplied across households. Instead we assume that the

Figure 4: Mean MPC of households close to adjustment point

disutility of labour supply is separable from consumption and thus varies with household wealth. This assumption decreases the prevalence of households with negative marginal propensities to consume, as households have a second margin which they can adjust as they anticipate buying or selling housing stock. However while the proportion of households with a negative marginal propensity to consume is smaller, the phenomenon still exists across a range of reasonable calibrations.

5   Debt and Consumption: Quantitative Results

In this section I compare the predictions from the general equilibrium model against empirical results derived from household surveys. To compute the key cross-section moments, I require household-level panel data on income, consumption and household debt. There are a variety of

techniques in the literature used to calculate parameters such as households' marginal propensity to consume, the most commonly used are the US Panel Study of Income Dynamics (PSID) and the US Consumer Expenditure Survey. However since the Consumer Expenditure Survey does not follow households who move address, it excludes those who buy or sell their home - the main population of interest. I thus use the PSID to verify the predictions.

5.1 Measuring Marginal Propensities to Consume

The first exercise I conduct is to compare the predicted marginal propensity to consume of the model with the data. The model generates three broad predictions concerning the dynamics of non-durable consumption: (i) that households close to their borrowing constraint have a high marginal propensity to consume; (ii) that households who have close to zero liquid wealth have a high marginal propensity to consume; and (iii) that households who are close to adjusting their stock of housing have a low, or even negative, marginal propensity to consume. As there is already considerable empirical evidence supporting the first two predictions as outlined in my literature review, I will focus on the more novel third prediction.

To estimate marginal propensities to consume I follow the identification strategy of Blundell, Pistaferri and Preston 2008, (BPP) and popularised by Kaplan and Violante (2010). BPP show that by assuming household income is governed by a process with a permanent and an i.i.d. component, then given appropriate theoretical restrictions, an estimator for households' marginal propensities to consume can be calculated with panel data on consumption and income.[12]

The estimator for the transmission of transitory income shocks to consumption is given by

$$\widehat{MPC} = \frac{\text{cov}\left(\Delta c_{it}, \Delta y_{i,t+1}\right)}{\text{var}\left(\Delta y_{i,t}, \Delta y_{i,t+1}\right)} \tag{24}$$

which can be estimated using an instrumental variable regression, where $\Delta c_{i,t}$ is regressed on $\Delta y_{i,t}$, instrumented by $\Delta y_{i,t+1}$ with panel data of at least three periods.

Kaplan and Violante (2010) show that this estimator is highly robust at identifying the true marginal propensity to consume in a wide variety of models, including models in which there is a forecastable component of future income. Berger et al. (2015) further show that this method

---

[12]Two assumptions are required for the estimator to be consistent. The first is that households have no information about future shocks, the second is that consumption is memory-less and does not vary in response to the lags of the transitory shock.

remains robust in the presence of a housing market with transaction costs and the option to rent. I replicate the exercises conducted by Kaplan and Violante (2010) and show that the true marginal propensity to consume is reliably measured in my model of jump-diffusion income processes.

5.1.1 PSID Data

While the PSID started collecting data in 1968, it only started surveying non-food consumption in 1999. Although it is possible to impute aggregate consumption by inverting the demand for food using data from the Consumer Expenditure Survey Blundell et al. (2004), this measure can have high levels of measurement error as shown by Attanasio and Weber (1995). Accordingly, I calculate a broader measure of consumption which is measured bi-annually from 1999 to 2015. In this sub-sample, the PSID surveys over 70 per cent of all consumption items available in the Consumer Expenditure Survey and includes expenditure on food, utilities, gasoline, car repairs, public transportation, childcare, education and medical services. Following Kaplan and Violante (2010), I first purge the data of non-model features by regressing log consumption and log income on year & cohort dummies and a range of demographic variables.

5.1.2 Measuring Expected Changes in Debt

To verify the model's predictions on the heterogeneity of households' marginal propensities to consume, I need to identify households who plan on adjusting their housing stock in the near future. Using actual changes in household balance sheets may bias our results due to unexpected shocks which affect both a household's consumption and their decision to take on debt.[13]

To control for this potential bias I use households' surveyed expectation that they will move in the next couple of years. Specifically the PSID asks households:

A51. Do you think you might move in the next couple of years?

If a household replies in the affirmative, this question is followed with:

A52. Would you say you definitely will move, probably will move, or are you more uncertain?

---

[13]For example an unexpected persistent, positive income shock may lead households to increase both consumption and their willingness to hold debt.

This variable, when lagged, is orthogonal to contemporaneous shocks and highly correlated with future decisions to buy new houses as shown by Table 5.

Table 5: PSID Surveyed Expectations

| | N | Share of households that: | | | |
		subsequently move	move into a larger house	move into a similar sized house	move into a smaller house
Do not expect to move	47,305	0.17	0.80	0.04	0.16
Expectation - uncertain	4,988	0.34	0.84	0.03	0.13
Expectation - probably	8,758	0.50	0.86	0.02	0.12
Expectation - definitely	13,123	0.72	0.87	0.03	0.10

5.1.3 Results

In order to compute the marginal propensities to consume, I first group households into buckets. I try three different strategies to identify households who anticipate buying or selling housing stock. First, I group households by their self-reported expectations in the previous survey. The second approach is to estimate a logit model of the probability a household will move, regressed on household lagged expectations, demographic variables and lags of family income and employment. The equation provides us with a continuous measure of the likelihood that a household will move. Households are then grouped into quartiles based on this estimated likelihood of moving. Although the majority of households who adjust their stock of housing increase their stock, both of these strategies combine households regardless of whether they are net buyers or net sellers. To separate out the minority of households who move into smaller houses, my third approach groups households by whether they actually did move in the subsequent period, and whether the house they moved into had a smaller, unchanged or larger number of rooms.

Figure 5 shows the estimated marginal propensities to consumes of households according to their expectations about the likelihood of moving in the next couple of years.

These results show that households with a higher expectation of moving house consistently have a lower marginal propensity to consume. Indeed for households that report that they "Definitely expect to move" in subsequent periods or which are in the top quartile of the estimated equation, the estimator is statistically below zero as predicted by the model. Furthermore when households are grouped by their actual change in housing tenure, I show that this effect is stronger
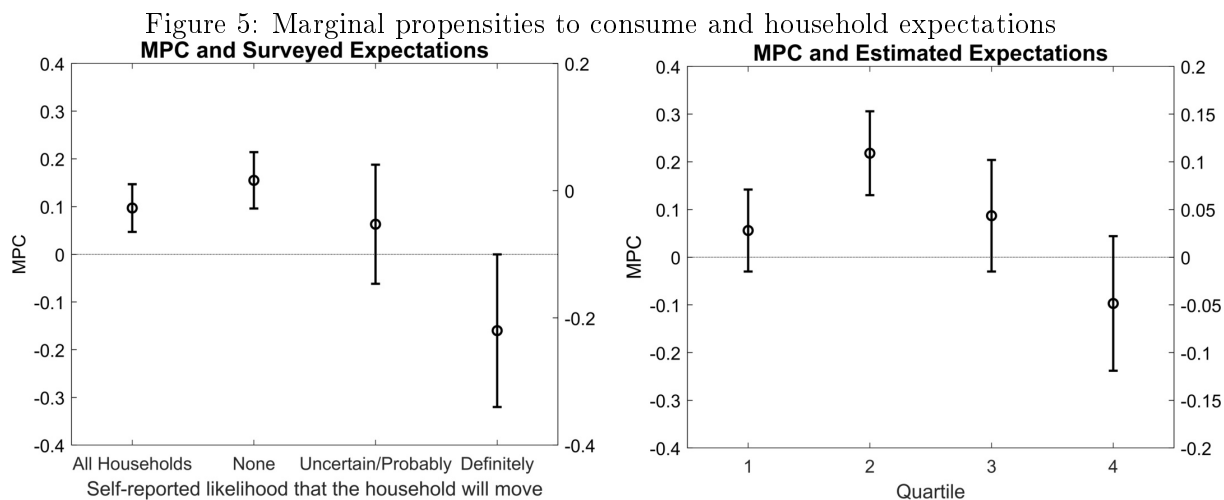
Figure 5: Marginal propensities to consume and household expectations

Note. This figure shows the estimated marginal propensity to consume with households grouped by their surveyed expectation of moving in the next two years (left) or by the likelihood that they will move estimated using a logit model regressing households' decisions to move on lagged expectations and demographic variables.

for those households who move into larger houses.

This result can be explained down the covariances which compromise the BPP estimator. The correlation between Δy_{t+1} and Δc_t is negative for the population as a whole, but is positive for the subset of households who report that they have a high probability of moving. This change in sign occurs because households who expect to move, but subsequently receive a negative income shock, are more likely to delay or cancel their planned purchase of housing. By choosing not to adjust their stock of housing they remain free from mortgage debt and subject to a lower interest rate, which results in a higher level of consumption.

While the standard errors around the estimates produced with this method are large, the covariance between anticipated increases in debt and households' marginal propensities to consume, and the negative result for those who are highly likely to buy additional housing stock are consistent with the results of the model.

The identification of households with a negative marginal propensities to consume is a novel result. The closest related work is Martin (2003) who shows a similar effect in a partial equilibrium model of the housing market with transaction costs. There are several possible explanations for why result has not occurred in the previous literature. First, the share of households who adjust their stock of housing in any given period is quite small. As empirical studies usually compare households across quantiles, unless the analysis deliberately isolates

Figure 6: MPC and changes in household tenure
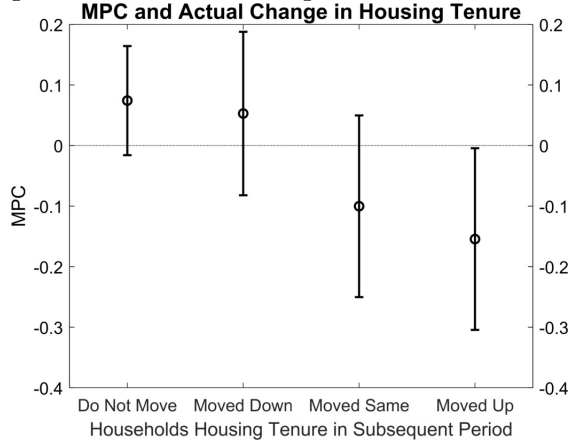
Note. This figure shows the estimated marginal propensity to consume with households grouped whether they actually moved in the subsequent two years and whether they moved into a house with less, the same or more rooms.

such households it is likely that they will be subsumed by the majority of households that do not adjust their stock of housing and have a conventionally non-negative marginal propensities to consume. Secondly, because the optimal adjustment point varies with income, even analyses using administrative data such as Fagereng et al. (2016), which is able to partition households by their balance sheets in finer gradients, may fail to find this effect. Finally, while the finding that some households have a negative marginal propensity to consume is new to the literature, these results are consistent with the broader conclusions on the negative correlation between households' marginal propensity to consume and their ownership of liquid assets.

5.2   Consumption Growth Regressions

The second exercise I conduct is to estimate how anticipated changes in debt affect the level of consumption of a household. To estimate this effect I adapt the model used by Dynan (2012) to estimate how changes in debt affect the growth in consumption,

$$\Delta \log c_{i,t} = \beta_0 + \beta_r r_{t-1} + \beta_y \Delta \log y_{i,t} + \beta_{Debt} \Delta D_{i,t} + \beta_W \Delta \log NW_{i,t} + \boldsymbol{\beta} \boldsymbol{X}_{i,t} + \varepsilon_{i,t},$$

where $c_{i,t}$ is the log of real consumption, r_t is the real interest rate, $y_{i,t}$ is the household's real income, $D_{i,t}$ is the debt-to-asset ratio of each household and $\boldsymbol{X}_{i,t}$ is a set of demographic

27

variables. Note that since the PSID is a bi-annual survey, I define Δ such that $\Delta x_t = x_t - x_{t-2}$. I use this model to examine (i) how anticipated changes in household debt and housing stock affect non-durable consumption and (ii) how consumption responds to changes in households' expectations of when they will adjust their stock of housing.

In order to identify the impact of expected debt on household consumption, I instrument the debt-to-asset ratio with households' lagged expectation of the probability that they will move along with lags for income, consumption and demographic variables. I extend Hansen and Singleton (1982) to show that if the set of instruments used to identify the endogenous variable contain only information from time $t-2$ then an instrumental variable regression will produce a consistent estimator of the expected change in debt, excluding any unanticipated changes. This result is driven by the assumption of rational expectations, in which i.i.d deviations from prior expectations are orthogonal to information available at time $t-2$. This extension is outlined in Appendix C. By using the household's self reported expectation of moving as an instrument for changes in debt, I produce a consistent estimator of the effect of an expected change in debt on household consumption. It should be noted that while this estimator is consistent, it is only unbiased when the marginal effect of an expected increase in debt is equal to that of an unexpected change. However, simulating the model under a wide range of parameters suggests that this bias will be quantitatively small given our sample size.

Table 6: Regression results

	(1)	(2)	(3)	(4)	(5)
				1999-2007	OLS
ΔDebt-Asset$_{i,t}$	-0.043 *	-0.043 *	-0.046 *	-0.051 *	0.011
ΔFamily Size$_{i,t}$	0.015	0.015	0.012	0.020	0.079***
r_{t-2}	-0.09***	-0.09***	-0.09***	0.00	-0.04*
$\Delta \log y_{i,t}$		-0.005	-0.005	-0.005	0.013***
$\Delta \log W_{i,t}$			-0.002	-0.005*	0.000
Observations	40,965	38,921	37,560	22,601	37,560

Source: Authors regressions of nested specifications of the model $\Delta \log c_{i,t} = \beta_0 + \beta_r r_{t-1} + \beta_y \Delta \log y_{i,t} + \beta_{Debt} \Delta D_{i,t} + \beta_W \Delta \log NW_{i,t} + \boldsymbol{\beta} \boldsymbol{X}_{i,t} + \varepsilon_{i,t}$. The dependent variable is the change in log of real consumption. All regressions instrument contemporaneous variables using lags of endogenous variables and households expectation of moving, unless specifically stated otherwise. Sample includes households with consecutive, complete set of interviews from 1999 to 2015. Also included, but not reported, are a constant, controls for the age and education of the household head. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

These results show that anticipated increases in household debt have a significant impact on household consumption across a range of specifications, even when the increases are anticipated

years in advance.

To compare these estimates to the predictions of the model I calculate the mean change in non-durable consumption that occurs when households adjust their stock of housing, instrumenting again with household expectations, in Table 7. My estimates suggest that households that buy their first house decrease their consumption by approximately 10 per cent, relative to households who choose to keep renting. This figure increases to 12 per cent if we only include households that take out a mortgage in the process. By contrast, households who buy additional housing stock (as proxied by an increase in the number of rooms when home owners move) decrease their non-durable consumption by approximately 7 per cent. The results show that even households who anticipate their adjustment of a housing stock decrease their non-durable consumption in response to the purchase.

These estimates are somewhat larger then the simulated outcome of the model, which predicts a fall in consumption of approximately 4 per cent upon upgrading housing stock. Part of this difference is due to unanticipated income shocks that exist within the model. A household that is subject to a positive income shock will simultaneously increase their consumption of both housing and non-durable goods. This will bias the simulated model results towards zero.

Table 7: Alternate regressions

	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)
ΔDebt-Asset$_{i,t}$	-0.04 *	-0.05 *						
ΔMortgage$_{i,t}$			-0.13***	-0.11***				
ΔFirst Home Owner$_{i,t}$					-0.11***	-0.10***		
ΔBuy additional stock$_{i,t}$							-0.06*	-0.07*
ΔFamily Size$_{i,t}$	0.02	0.01	0.02	0.02	0.02 *	0.02 *	0.01	0.01
r_{t-2}	-0.09***	-0.09***	-0.05***	-0.05***	-0.05***	-0.06***	-0.01***	-0.01***
$\Delta \log y_{i,t}$	-0.01	-0.01	0.03 *	0.03 *	0.03 *	0.03 *	0.02	0.020
$\Delta \log W_{i,t}$		-0.00		-0.03		-0.03		-0.00
Observations	37,560	37,560	37,560	37,560	37,560	37,560	37,560	37,560

Source: Authors regressions of the baseline model with alternative measures of debt and home ownership. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

These results, using a more comprehensive measure of consumption, confirm the findings from Martin (2003). Namely, that households which choose to purchase additional housing stock decrease their non-durable consumption. This result accords with the predictions of the model and highlights the importance of transaction costs and financial frictions when modeling

a household's consumption dynamics.

Table 8: Changes in household expectations

		Period t			
		No expectation	Uncertain	Probably	Definitely
	No expectation	-	−0.011	−0.011	−0.073***
Period $t-2$	Uncertain	−0.015	−0.031	−0.010	−0.063***
	Probably	0.007	−0.017	−0.027*	−0.048***
	Definitely	0.020*	−0.065	−0.029	−0.057***

Estimated impact of a change in household expectations on household consumption.
Source: Authors regressions of nested specifications of the model $\Delta \log c_{i,t} = \beta_0 + \beta_r r_{t-1} + \beta_y \Delta \log y_{i,t} + \boldsymbol{\beta_{Exp}} \boldsymbol{\Delta E_{i,t}} +$
$\beta_W \Delta \log NW_{i,t} + \boldsymbol{\beta X}_{i,t} + \varepsilon_{i,t}$, where $\boldsymbol{\Delta E_{i,t}}$ is a vector of dummy variables which represent all potential transitions
in household expectations across two periods. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Finally, I also find evidence that consumption of non-durable goods decreases as household expectations of moving rise, even when their stock of housing remains constant. Table 8 shows the estimated elasticity of changes in household expectations on consumption. I find that households who definitely expect to move have lower levels of consumption, and that the fall is larger for households that report a bigger change in expectations. By contrast there is evidence that households who report a lower probability that they will move have a higher level of consumption. These results are consistent with the predictions from my model in which households decrease their level of non-durable consumption in anticipation of making a leveraged purchase of housing in the near future.

These results are in line with Benmelech et al. (2017), who find that non-durable consumption falls after the purchase of a house. However my results indicate that the decrease is significantly larger. this difference is driven by the use of household expectations as a instrumental variable. The OLS regressions results are much more similar in magnitude. The instrumental variable approach removes the impact of contemporaneous shocks that may impact both non-durable consumption and the decision to purchase housing. However it should be noted that the PSID does not cover the consumption of durable goods. Benmelech et al. (2017) show that the consumption of house-relate durable goods, such as furniture and electrical goods, is higher for households who have recently purchased a house. The exclusion of these measures likely biases my results downwards, in comparison to a more comprehensive measure of consumption.

6   Fiscal Transfers and the Housing Market

Subsidies for home ownership are a common feature of housing markets (Andrews et al., 2011). These policies range from subsidised access to credit, tax preferences for household debt and outright transfers to home buyers. While many of these subsidies are structural in nature, a common component of the fiscal response to the global financial crisis was to expand these programs to help stabilise housing markets and stimulate demand within the economy (Table 9).

The results from the previous sections raise the question of whether a transfer payment or tax credit to these almost home-owners may decrease aggregate consumption in general equilibrium and thus be contractionary in nature. This is an important policy question as tax credits to households saving for their first deposit were enacted during the financial crisis in a number of countries. While these programs were primarily designed to help boost the housing and credit markets, it was also assumed that they would help boost the broader economy and increase output (Dynan et al., 2013).

Table 9: Selected fiscal transfers targeted at first home buyers

Country	Size (% of median home price)	Dates	Notes
US	$7,500-8,000 (3.5%)	April 2008 - June 2010	Initially an interest free loan, it was converted into an outright grant for first home buyers.
UK	Discount of up to $140,000 for public housing tenants to buy their home (35-70% depending on length of tenancy)	2012 - Present	"Right to Buy" discounts were introduced in 1980, but greatly expanded in 2012.
Australia	$11,200 (4 %)	October 2008 - September 2009	This is in addition to $6,000 permanently available.
Canada	$700 (0.3%)	January 2009 - Present	

I consider two policy scenarios. The first is an unconditional transfer of money to renters who are on the margin of purchasing a house. The second is a conditional tax credit which is available to all renters who purchase additional housing stock within a one year window.

6.1 Unconditional Transfers

To estimate the impact of these programs, I model the impact of a $8,000 debt-funded transfer to renters who are sub-marginal first home buyers. I model this transfer as a one-off increase in liquid wealth to households who are currently renting and are on the margin of becoming a home owner. The transfer is funded by the government through an increase in debt, while tax rates remain unchanged. I then solve for the impact of this transfer in equilibrium allowing prices to adjust to clear markets.

Figure 7: Impact of a unconditional transfer
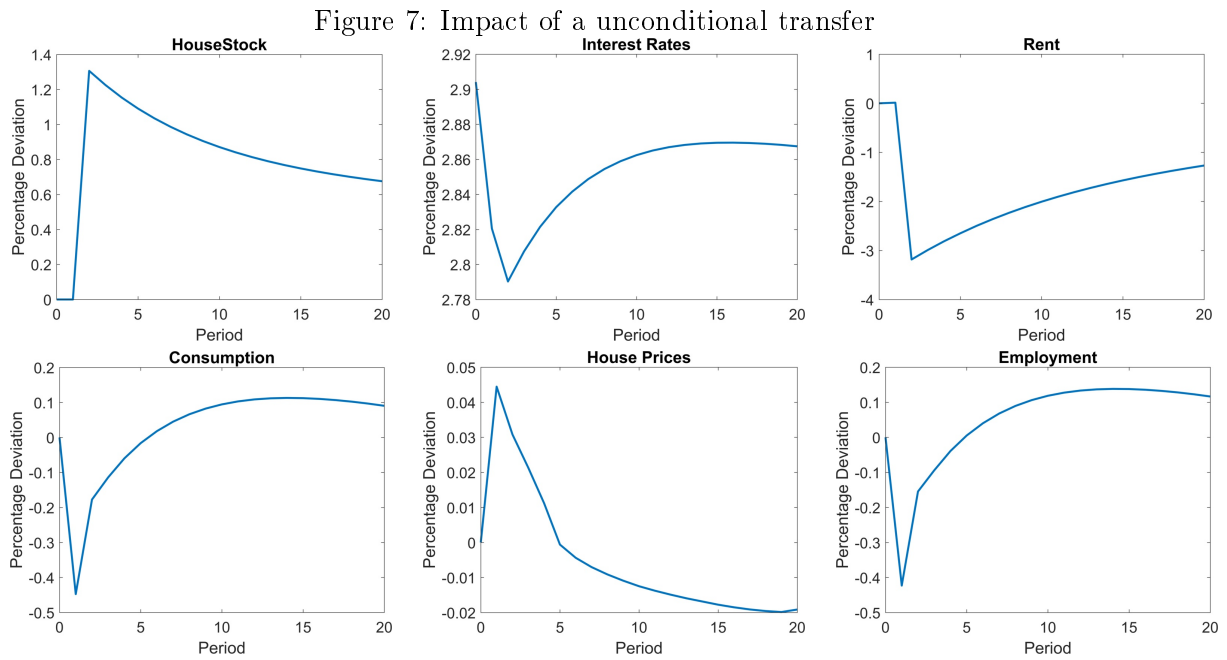
Note. These impulse response functions show the impact of an unconditional $8,000 transfer to renters

I solve for the equilibrium path for the model as it returns to steady state after the initial transfer. Figure 7 shows the impact from this policy on the housing market and the broader economy. Unsurprisingly, the transfer increases both prices and the stock of housing owned by households, as the transfer to renters enables them to purchase more housing and to purchase housing sooner than they otherwise would have. This leads to an increase in household leverage within the model which results in households decreasing their consumption via the mechanism that I explored than when prices and wages are held constant.

6.2 Conditional Transfers

An alternative policy design is to offer tax credits when renters purchase a house. I model this scenario by announcing a policy that will give renters a $8,000 payment if they purchase a house in the next four quarters. This policy differs from a straight fiscal transfer in two ways. First, by applying the tax credit over a longer period, it allows more renters to take advantage of magnifying its impact on the goods and housing market. The second is that the tax credit provides the opportunity for renters to adjust their behavior in anticipation of taking advantage of the credit before it expires.

Figure 8: Impact of a conditional transfer
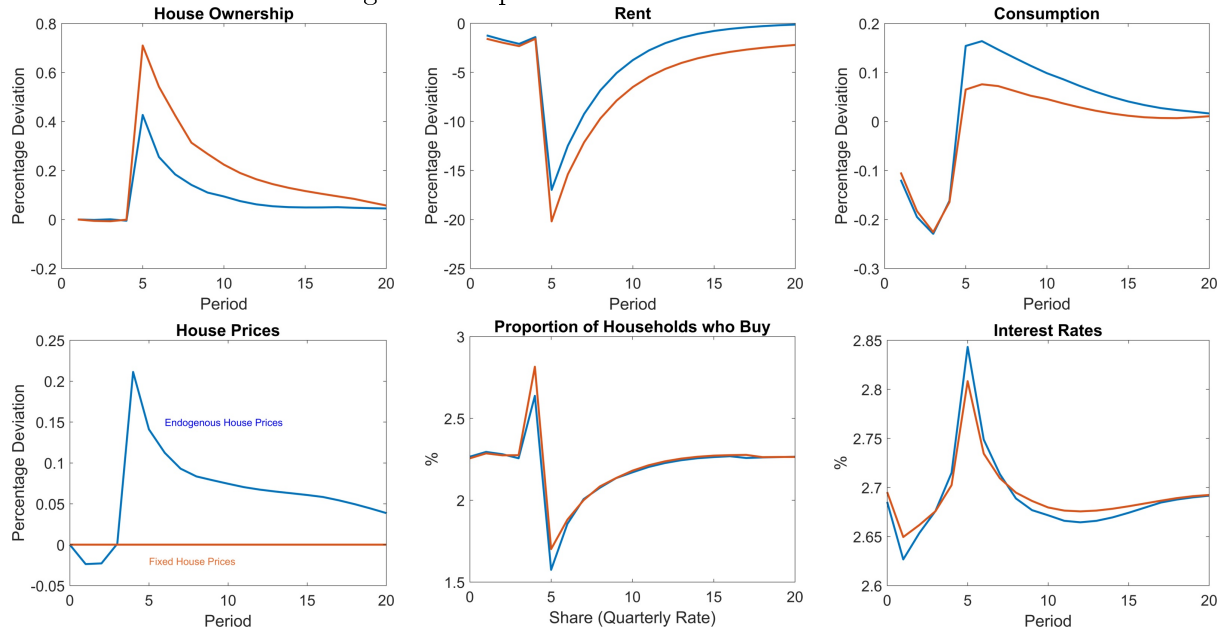
Note. These impulse response functions show the impact of an $8,000 transfer to renters conditional on buying housing stock within the first year. The results are presented under the baseline assumption that the price of housing adjusts to clear the market, and the alternative assumption that the price is held constant.

If the tax policy is open to all renters who buy a house within a set period, in this case one year, the general equilibrium impact of the policy has two distinct phases. From the time the tax credit is announced, renters decrease their consumption of both non-durable goods and rental stock in order to increase their stock of savings to a level high enough to enable them to buy a house while the credit is available. This temporary incentive to purchase a house thus results in a fall in consumption while households try to increase their stock of savings. When the tax credit is due to expire, renters' incentive to buy a house rather than delay and continue saving

spikes, and demand for housing increases dramatically. This leads to an increase in purchases, house prices and interest rates on debt.

While home owners who are not able to access the tax credit directly, they are affected indirectly by the change in prices within the economy. The decline in house prices and wages and the increase in interest rates while the tax credit is available causes home owners to lower their non-durable consumption. Anticipating the increase in house prices that occurs just before the tax credit expires, home owners are incentivised to bring forward (delay) any increase (decrease) in their stock of housing, however this effect is quantitatively small. After the tax credit expires home owners consumption increases as the rise in house prices increases their wealth and the lower interest rate decreases their desire to save.

Although liquid assets are also in zero net-supply, the variance of the distribution varies within the economy. The tax credit encourages renters to take on high levels of leverage in order to purchase housing which increases the variance of the distribution.

After the tax credit has expired the economy slowly returns to equilibrium. The overall result is to effect a shift in housing from the real estate sector to owners, due to the incentive provided by the temporary tax credit which pulls forward housing purchases from future periods. The incentive to pull forward the purchase of housing has the opposite effect on non-durable consumption, which is pushed back as households prioritise their investment in housing. Thus despite the stated intention of policy makers to stimulate the broader economy, I find that this tax credit has a contractionary impact on aggregate consumption and output.

Previous empirical studies of these programs (Dynan et al., 2013) found that the tax credit affected the price of housing, they did not find a significant positive impact on the broader economy. These results suggest that one possible resolution to this puzzle is the tax credit merely encourages households to substitute consumption across time and goods, pulling forward the purchase of housing and pushing back the consumption of non-durable goods. This finding is consistent with Mian and Sufi (2012), who evaluated other temporary tax credits such as the "cash for clunkers" program, which offered tax credits for owners trading in fuel inefficient cars, and show that it had a limited effect on overall demand.

7 Conclusion

This paper extends the two-asset Bewley model to allow for endogenous borrowing constraints determined by households' ownership of illiquid assets. I establish that the presence of financial frictions and transaction costs in the housing market generate important distortions in households' consumption of non-durable goods. In particular, I show that these frictions create an anticipatory saving effect. As the probability that a household will choose to adjust their stock of housing rises, their level of consumption will fall. This occurs because households anticipate that they will face higher interest rates and be credit constrained after making a debt-financed purchase of housing.

My results suggest that this effect is quantitatively meaningful. For households close to their optimal point of adjustment, an increase in liquid assets leads to a decrease in consumption as the anticipatory saving effect outweighs the wealth and liquidity effects. I provide empirical evidence for these results using micro-data from the PSID to show that a household's marginal propensity to consume is negatively correlated with their expectation of moving. I find that households which definitely expect that they will move have a negative marginal propensity to consume.

This result has important implications for the use of fiscal incentives in the housing market as a counter-cyclical policy tool. My general equilibrium model shows that tax credits for first home buyers decrease aggregate consumption due to an increase in mortgage debt and a higher incentive to save. These tax credits do boost prices and sales in the housing market, though much of this is the result of demand being pulled forward. While my analysis indicates that such policies lower output and employment, the financial system in my model is relatively simple. In the context of a bursting house price bubble and a financial crisis, a policy that helps to stabilise house prices may be welfare-improving despite the decrease in consumption due to higher household leverage.

More broadly, these results suggest that the considerable evidence of heterogeneity in households' marginal propensities to consume has significant consequences for the optimal design of macroeconomic policy. Beyond the implications for fiscal transfers my work on the relationship between debt and consumption raises new questions, for example how this mechanism may affect the transmission of monetary policy, which I leave as an avenue for future work.

References

Achdou, Y., Han, J., Lasry, J.M., Lions, P.L., Moll, B., 2017. Income and wealth distribution in macroeconomics: A continuous-time approach. Technical Report. National Bureau of Economic Research.

Ahn, S., Kaplan, G., Moll, B., Winberry, T., Wolf, C., 2017. When inequality matters for macro and macro matters for inequality, in: NBER Macroeconomics Annual 2017, volume 32. University of Chicago Press.

Aiyagari, S.R., 1994. Uninsured idiosyncratic risk and aggregate saving. The Quarterly Journal of Economics 109, 659–684.

Andrews, D., Sanchez, A.C., Johansson, Å., 2011. Housing markets and structural policies in oecd countries. OECD Economic Department Working Papers , 0_1.

Attanasio, O.P., Bottazzi, R., Low, H.W., Nesheim, L., Wakefield, M., 2012. Modelling the demand for housing over the life cycle. Review of Economic Dynamics 15, 1–18.

Attanasio, O.P., Weber, G., 1995. Is consumption growth consistent with intertemporal optimization? evidence from the consumer expenditure survey. Journal of political Economy 103, 1121–1157.

Auclert, A., 2017. Monetary policy and the redistribution channel. Technical Report. National Bureau of Economic Research.

Benmelech, E., Guren, A., Melzer, B.T., 2017. Making the House a Home: The Stimulative Effect of Home Purchases on Consumption and Investment. Technical Report. National Bureau of Economic Research.

Berger, D., Guerrieri, V., Lorenzoni, G., Vavra, J., 2015. House prices and consumer spending. Technical Report. National Bureau of Economic Research.

Best, M.C., Kleven, H.J., 2013. Housing market responses to transaction taxes: Evidence from notches and stimulus in the uk. London School of Economics .

Bewley, T., 1983. A difficulty with the optimum quantity of money. Econometrica: Journal of the Econometric Society , 1485–1504.

Blundell, R., Pistaferri, L., Preston, I., 2004. Imputing consumption in the psid using food demand estimates from the cex .

Blundell, R., Pistaferri, L., Preston, I., 2008. Consumption inequality and partial insurance. The American Economic Review 98, 1887–1921.

Broda, C., Parker, J.A., 2014. The economic stimulus payments of 2008 and the aggregate demand for consumption. Journal of Monetary Economics 68, S20–S36.

Caballero, R.J., Farhi, E., 2014. On the role of safe asset shortages in secular stagnation. Secular stagnation: Facts, causes and cures , 111.

Campbell, J.Y., Cocco, J.F., 2007. How do house prices affect consumption? evidence from micro data. Journal of monetary Economics 54, 591–621.

Carroll, C.D., Dunn, W.E., 1997. Unemployment Expectations, Jumping (S,s) Triggers, and Household Balance Sheets. Working Paper 6081. National Bureau of Economic Research. URL: http://www.nber.org/papers/w6081, doi:10.3386/w6081.

Dynan, K., 2012. Is a household debt overhang holding back consumption? Brookings Papers on Economic Activity 2012, 299–362.

Dynan, K., Gayer, T., Plotkin, N., 2013. An evaluation of federal and state homebuyer tax incentives. Washington, DC: The Brookings Institution .

Fagereng, A., Holm, M.B., Natvik, G.J., 2016. Mpc heterogeneity and household balance sheets .

Favilukis, J., Ludvigson, S.C., Van Nieuwerburgh, S., 2017. The macroeconomic effects of housing wealth, housing finance, and limited risk sharing in general equilibrium. Journal of Political Economy 125, 140–223.

Fernandez-Villaverde, J., Krueger, D., 2011. Consumption and saving over the life cycle: How important are consumer durables? Macroeconomic dynamics 15, 725–770.

Garriga, C., Hedlund, A., 2016. Mortgage Debt, Consumption, and Illiquid Housing Markets in the Great Recession. Technical Report. Working Paper.

Greenwood, J., Hercowitz, Z., Huffman, G.W., 1988. Investment, capacity utilization, and the real business cycle. The American Economic Review , 402–417.

Greer, J.M., Parker, J.A., Souleles, N.S., 2006. Household expenditure and the income tax rebates of 2001. The American Economic Review 96, 1589–1610.

Guerrieri, L., Iacoviello, M., 2017. Collateral constraints and macroeconomic asymmetries. Journal of Monetary Economics 90, 28–49.

Guvenen, F., Karahan, F., Ozkan, S., Song, J., 2015. What do data on millions of US workers reveal about life-cycle earnings risk? Technical Report. National Bureau of Economic Research.

Hansen, L.P., Singleton, K.J., 1982. Generalized instrumental variables estimation of nonlinear rational expectations models. Econometrica: Journal of the Econometric Society , 1269–1286.

Huggett, M., 1993. The risk-free rate in heterogeneous-agent incomplete-insurance economies. Journal of economic Dynamics and Control 17, 953–969.

Iacoviello, M., Pavan, M., 2013. Housing and debt over the life cycle and over the business cycle. Journal of Monetary Economics 60, 221–238.

Justiniano, A., Primiceri, G.E., Tambalotti, A., 2015. Credit supply and the housing boom. Technical Report. National Bureau of Economic Research.

Kaplan, G., Moll, B., Violante, G.L., 2016. Monetary policy according to HANK. Technical Report. National Bureau of Economic Research.

Kaplan, G., Violante, G.L., 2010. How much consumption insurance beyond self-insurance? American Economic Journal: Macroeconomics 2, 53–87.

Kaplan, G., Violante, G.L., 2014. A model of the consumption response to fiscal stimulus payments. Econometrica 82, 1199–1239.

Kaplan, G., Violante, G.L., Weidner, J., 2014. The wealthy hand-to-mouth. Technical Report. National Bureau of Economic Research.

Krusell, P., Smith, Jr, A.A., 1998. Income and wealth heterogeneity in the macroeconomy. Journal of political Economy 106, 867–896.

Martin, R.F., 2003. Consumption, durable goods, and transaction costs .

Mian, A., Rao, K., Sufi, A., 2013. Household balance sheets, consumption, and the economic slump. The Quarterly Journal of Economics 128, 1687–1726.

Mian, A., Sufi, A., 2012. The effects of fiscal stimulus: Evidence from the 2009 cash for clunkers program. The Quarterly journal of economics 127, 1107–1142.

Mian, A., Sufi, A., 2014. What explains the 2007–2009 drop in employment? Econometrica 82, 2197–2223.

Misra, K., Surico, P., 2014. Consumption, income changes, and heterogeneity: Evidence from two fiscal stimulus programs. American Economic Journal: Macroeconomics 6, 84–106.

Ngai, L.R., Sheedy, K.D., et al., 2016. The Decision to Move House and Aggregate Housing-Market Dynamics. Technical Report.

Parker, J.A., Souleles, N.S., Johnson, D.S., McClelland, R., 2013. Consumer spending and the economic stimulus payments of 2008. The American Economic Review 103, 2530–2553.

Saiz, A., 2010. The geographic determinants of housing supply. The Quarterly Journal of Economics 125, 1253–1296.

Yang, F., 2009. Consumption over the life cycle: How different is housing? Review of Economic Dynamics 12, 423–443.

A Details on Model Solution

A.1 Model Approximation

I present the households' Hamilton-Jacobi-Bellman (HBJ) equation and Kolmogorov forward equation for the evolution of the distribution of households μ over their holdings of liquid assets, b, housing stock, a, and income, z. We first solve for the HJB equation assuming that households choose not to adjust their stock of housing. The stationary version of the households' HJB equation is given by

$$
\begin{aligned}
(\rho + \lambda) V^s (a, b, z) = & \max_{c, l, c^r, \alpha} u (c, h, l) + V_b^s \left[zw (1 - \tau) l + r (b) b - c - c^{rent} p^{rent} + T \right] \quad (25) \\
& + V_a^s (\delta^a a) + V_z^s (-\beta z) + \lambda \int_{-\infty}^{\infty} (V^s (a, b, x) - V^s (a, b, z)) \phi (x) \, dx,
\end{aligned}
$$

where the value function is determined by the current flow of utility and the expected change in state variables.

The optimal drift in liquid assets is equal to the level of household savings which is governed by the budget constraint. Since we assume that households do not buy or sell housing stock the optimal drift in housing stock is driven solely by depreciation. Finally, the expected change in household income is determined by the mean reversion and the expectation of idiosyncratic "jumps" which cause log-income to jump to a new level drawn from normal distribution of density ϕ with variance σ^2.

We assume that if a household chooses to adjust their stock of housing than they will choose the optimal portfolio of assets subject to their budget constraint

$$
V^a (a, b, z) = \max V^s (a', b', z) \tag{26}
$$

$$
\text{s.t.}
$$

$$
p^a a + b = p^a a' + b' - k (a, a') .
$$

The household re-allocates their wealth across the two assets subject to paying the cost of adjustment, $k (a, a')$.

The households' final value function is given by

$$V(a, b, z) = \max\left\{V^a(a, b, z,), V^s(a, b, z,)\right\},$$

such that households choose to "stay" or "adjust" depending on which value function is larger. The adjustment function can then be defined as

$$I^{adjust} = \begin{cases} 1 & \text{if } V(a, b, z) = V^a(a, b, z) \\ 0 & \text{if } V(a, b, z) = V^s(a, b, z), \end{cases} \tag{27}$$

and, conditional on paying the adjustment cost, the optimal portfolio, $a^*(a, b, z)$ and $b^*(a, b, z)$, is solved by maximising 26.

The households' optimal choice variables $\left(c_t, c_t^{rent}, l_t\right)$ are determined by this value function and the inverse of the utility function

$$c_{i,j,k} = \left(\frac{\partial u}{\partial c}\right)^{-1}\left[\left(\frac{\partial V(a, b, z,)}{\partial b}\right)\right] \tag{28}$$

$$c_{i,j,k}^{rent} = \left(\frac{\partial u}{\partial h}\right)^{-1}\left[\frac{p_t^r}{\theta^r}\left(\frac{\partial V(a, b, z,)}{\partial b}\right)\right] \tag{29}$$

$$l_{i,j,k} = \left(\frac{\partial u}{\partial l}\right)^{-1}\left[w_t\left(\frac{\partial V(a, b, z,)}{\partial b}\right)\right] \tag{30}$$

The evolution of the joint distribution of households is described by a Kolmogorov forward equation. I denote the density of households over the joint-distribution $g(a, b, z, t)$ corresponding to the density function $\mu_t(a, b, z)$. Furthermore I denote the optimal drift function for liquid assets as $s(a, b, z)$. Then the stationary density must satisfy the Kolmogorov forward equations

$$\begin{aligned} 0 = {}& -\partial_b\left(s(a, b, z)\, g(a, b, z)\right) \\ & -\partial_a\left(\delta^a g(a, b, z)\right) + \int I^{adjust}\delta(a - a^*)\,\delta(b - b^*)\ d\mu_t \\ & -\partial_z\left(-\beta z g(a, b, z)\right) - \alpha g(a, b, z) + \alpha\phi(z)\int_{-\infty}^{\infty} g(a, b, x)\ dx \\ & -\lambda g(a, b, z) + \lambda\delta(a - a_0)\,\delta(b - b_0)\,\delta(z - z_0), \end{aligned} \tag{31}$$

41

where δ is the Dirac delta function and (a_0, b_0, z_0) are the birth level of assets and income.

This equation equates inflows and outflows of households across each point in the joint-distribution. The first line measures the outflow as households increase (decrease) their stock of liquid assets through saving (dis-saving). The second line relates to changes in households' stock of housing via depreciation and the decision to buy or sell. The third line describes how exogenous changes in household productivity impact the distribution. Finally, the stochastic death and the birth of households with zero wealth is described in the last line.

A.2 Steady State

I solve (25) and (31) expanding on the finite difference method outlined by Achdou et al. (2017). To solve the HJB function for households who choose not to adjust their stock of housing, I approximate the function $V^s (a, b, z)$ on $I \times J \times K$ discrete equi-spaced points. I denote the grid points $b_i, i = 1, ... I$, $a_j, j = 1, ..., J$, $z_k, k = 1, ... K$, and define the discretised value function for households that do not adjust their stock of housing

$$v_{i,j,k} = V^s (a_j, b_i, z_k).$$

I approximate the derivatives for b with either a forward or backward difference approximation

$$
\begin{aligned}
\frac{\partial v_{i,j,k}}{\partial b} &\approx v_b^F = \frac{v_{i+1,j,k} - v_{i,j,k}}{\Delta b_i^+} \\
\frac{\partial v_{i,j,k}}{\partial b} &\approx v_b^B = \frac{v_{i,j,k} - v_{i-1,j,k}}{\Delta b_i^-},
\end{aligned}
\tag{32}
$$

and use an upwind method to choose when to use each approximation. I use the forward (backward) approximation whenever the drift of b is positive (negative). Since the drift terms for housing and income are always negative (due to depreciation and mean reversion) I use the backward difference approximation universally for these variables. As outlined by Achdou et al. (2017), this approximation allows us to solve for the discretised value function using a process of iteration.

$$\frac{v_{i,j,k}^{n+1} - v_{i,j,k}^n}{\Delta^{step}} + (\rho + \lambda)\, v_{i,j,k}^{n+1} = u\left(c_{i,j,k}^n, h_{i,j,k}^n, l_{i,j,k}^n\right)$$
$$+ v_b^n \left[z_t w_t (1 - \tau)\, l_{i,j,k}^n + r\,(b_i)\, b_i - c_{i,j,k}^n - c_{i,j,k}^{rent,n} p_t^{rent} + T_t \right]$$
$$+ v_a^n\,(\delta^a a_j) + v_z^n\,(-\beta z_k) + \lambda \int_{-\infty}^{\infty} \left(v_{i,j,x}^n - v_{i,j,k}^n \right) \phi\,(x)\, dx,$$

which can in turn be written as

$$\left[\left(\frac{1}{\Delta^{step}} + \rho + \lambda \right) \boldsymbol{I} - \boldsymbol{A} \right] v_{i,j,k}^{n+1} = u\left(c_{i,j,k}^n, h_{i,j,k}^n, l_{i,j,k}^n \right) + \frac{1}{\Delta^{step}} v_{i,j,k}^n, \tag{33}$$

where $\left(c_{i,j,k}^n, h_{i,j,k}^n, l_{i,j,k}^n \right)$ are the choice variables optimised by the households under the value function $v_{i,j,k}^n$ and $\boldsymbol{A^n}$ is a $(I \times J \times K)^2$ square matrix which summarises the exogenous productivity process, depreciation of housing stock and optimal household saving.

The HJB equation is then solved for with the following algorithm.

1. Start with a guess for the value function $v_{i,j,k}^0$

2. Compute the derivatives of $v_{i,j,k}^n$ using (32) and the upwind approximation.

3. Compute the optimal household choices $\left(c_{i,j,k}^n, h_{i,j,k}^n, l_{i,j,k}^n \right)$ using the first order conditions X and the approximated derivatives of $v_{i,j,k}^n$

4. Find $v_{i,j,k}^{n+\frac{1}{2}}$ from (33)

5. Compute the value function for households conditional of moving with

$$\left(v_{i,j,k}^* \right)^{n+\frac{1}{2}} = \max_{a',b'} v_{i,j,k}^{n+\frac{1}{2}} \;\text{s.t.}\; p^a a + b = p^a a' + b' - k\,(a, a').$$

6. Given $v_{i,j,k}^{n+\frac{1}{2}}$ and $\left(v_{i,j,k}^* \right)^{n+\frac{1}{2}}$ calculate

$$v_{i,j,k}^{n+1} = \max \left\{ v_{i,j,k}^{n+\frac{1}{2}}, \left(v_{i,j,k}^* \right)^{n+\frac{1}{2}} \right\}$$

7. Repeat steps 2 through 6 until the difference between $v_{i,j,k}^{n+1}$ and $v_{i,j,k}^n$ is below a given threshold.

Once the discretised value function has been calculated we can solve the Kolmogorov Forward equation. In the absence of adjustments in the stock of housing, the joint distribution can be conveniently be calculated by the linear system

$$0 = \boldsymbol{A}'g,$$

where $\boldsymbol{A}$ is the transpose of the transition matrix in (33) as shown by Achdou et al. (2017), with a minor adjustment for the stochastic birth and death process.[14]

To account for adjustments in the stock of housing we need to modify this transition matrix to

$$\boldsymbol{C}_{m,n} = \begin{cases} 0, & \text{if } I^{adjust}(n) = 1 \\ \boldsymbol{A}_{m,n} + \sum_{I^{adjust}(n)=1} \boldsymbol{A}_{m,n} & \text{if } I^{adjust}(n) = 0, \quad m'(m) = n \\ \boldsymbol{A}_{m,n} & \text{if } I^{adjust}(n) = 0, \quad m'(m) \neq n, \end{cases} \quad (34)$$

where m' is the optimal adjustment target conditional on a household deciding to change its stock of housing. I^{adjust} is a vector which indicates if it is optimal for a household to adjust.

The first row indicates that any point n that lies in the adjustment region will have a density of 0. Instead, the second row says that any households who would otherwise transition into the adjustment region instead move to the optimal adjustment target. Finally, any point that is not in the adjustment region, nor the optimal target for a point in the region, remains the same as the original transition matrix.[15] The density is thus given by

$$0 = \boldsymbol{C}'g.$$

where $g(a, b, z)$ is normalised such that

$$1 = \int_0^{a_J} \int_{b_1}^{b_I} \int_{z_1}^{z_K} g(a, b, z) \, da \, db \, dz.$$

Once we have solved for the density of households, $g(a, b, z)$, we can easily solve for the

---

[14]Specifically, I add, λ to each element of the row associated with the birth position (a_0, b_0, z_0) and subject λ from each diagonal element of $\boldsymbol{A}$.

[15]To ensure the matrix $\boldsymbol{C}$ is not singular we also set $\boldsymbol{C}_{m,m} = -\varepsilon$ for some small $\varepsilon > 0$, for all $I^{adjust}(m) = 1$. This ensures that all points in the adjustment region have a defined density of zero.

aggregate demand for liquid, B_t, and illiquid, A_t, assets

$$B_t = \int_0^{a_J} \int_{b_1}^{b_I} \int_{z_1}^{z_K} bg\,(a,b,z)\ da\,db\,dz$$

$$A_t = \int_0^{a_J} \int_{b_1}^{b_I} \int_{z_1}^{z_K} ag\,(a,b,z)\ da\,db\,dz.$$

A.3   Transition Dynamics

Outside of the steady state, when changes in prices, wages or interest rates vary over the time, the time-dependent HJB equation is required

$$
\begin{aligned}
(\rho + \lambda)\,V^s\,(a,b,z,t) \;=\; & \max_{c,l,c^r,\alpha} u\,(c,h,l) + V_b^s\left[z_t w_t\,(1-\tau)\,l + r_t\,(b)\,b - c - c^{rent} p_t^{rent} + T_t\right] \\
& + V_a^s\,(\delta^a a) + V_z^s\,(-\beta z) + \lambda \int_{-\infty}^{\infty} (V^s\,(a,b,x) - V^s\,(a,b,z))\,\phi\,(x)\ dx \\
& + V_t^s\,(a,b,z,t)\,,
\end{aligned}
$$

where the last term accounts for changes in the value function over time. Solving this function over the same discrete grid as in the steady state results in the approximation

$$
\begin{aligned}
(\rho + \lambda)\,v_{i,j,k}^n \;=\; & u\left(c_{i,j,k}^n, h_{i,j,k}^n, l_{i,j,k}^n\right) + v_b^n\left[z_t w_t\,(1-\tau)\,l_{i,j,k}^n + r\,(b_i)\,b_i - c_{i,j,k}^n - c_{i,j,k}^{rent,n} p_t^{rent} + T_t\right] \\
& + v_a^n\,(\delta^a a_j) + v_z^n\,(-\beta z_k) + \lambda \int_{-\infty}^{\infty} \left(v_{i,j,x}^n - v_{i,j,k}^n\right)\,\phi\,(x)\ dx \qquad (35) \\
& + \frac{v_{i,j,k}^{n+1} - v_{i,j,k}^n}{\Delta t}\,,
\end{aligned}
$$

where $v_{i,j,k}^n$ is now short-hand for $v\,(a_j, b_i, z_k, t^n)$ and the terminal condition $v_{i,j,k}^N$ is equal to the steady state solution $v_{i,j,k}$. This can be written as

$$\left(\rho + \lambda - \left[\boldsymbol{A}^{n+1}\right] + \frac{1}{\Delta t}\right) v_{i,j,k}^n = u\left(c_{i,j,k}^{n+1}, h_{i,j,k}^{n+1}, l_{i,j,k}^{n+1}\right) + \frac{1}{\Delta t}\left(v_{i,j,k}^{n+1}\right), \qquad (36)$$

which yields a solution for $v_{i,j,k}^n$ given $v_{i,j,k}^{n+1}$.

Similarly the joint distribution may vary over time, thus outside of steady state the time-

dependent Kolmogorov equation is

$$
\begin{aligned}
\partial_t g\left(a, b, z, t\right) = & - \partial_b\left(s\left(a, b, z\right) g\left(a, b, z\right)\right) \\
& - \partial_a\left(\delta^a g\left(a, b, z\right)\right) + \int I^{adjust}\delta\left(a - a^*\right)\delta\left(b - b^*\right)\, d\mu_t \\
& - \partial_z\left(-\beta z g\left(a, b, z\right)\right) - \alpha g\left(a, b, z\right) + \alpha\phi\left(z\right)\int_{-\infty}^{\infty} g\left(a, b, x\right)\, dx \qquad (37) \\
& - \lambda g\left(a, b, z\right) + \lambda\delta\left(a - a_0\right)\delta\left(b - b_0\right)\delta\left(z - z_0\right).
\end{aligned}
$$

Solving this equation over the same discrete grid requires a two-step n

$$
g^{n+1} = \left(\boldsymbol{I} - \Delta t(\boldsymbol{C^n})'\right)^{-1} g^n, \qquad (38)
$$

where $\boldsymbol{C}^n$ is the transition matrix from (36) adjusted for households who choose to adjust their stock of housing.

Once we have solved for the discreteised steady state value function, $v\left(a, b, z\right)$, and density $g\left(a, b, z\right)$ it is relatively straightforward to solve for transition dynamics in response to a "MIT" shock - an unanticipated shock followed by a deterministic transition back to the (potentially new) steady state. To solve for the transition path we use the following algorithm.

1. Guess a path for wages, w_t, interest rates, r_t^b, and house prices, p_t^a.

2. Given the path of wages, interest rates and prices solve the HJB backwards in time, using $v_{i,j,k}$ as the terminal condition.

3. Using this time path for the value function, calculate the optimal rules for household consumption and housing stock adjustment for each time period.

4. Given the optimal saving and adjustment rules, solve the time-dependent Kolmogorov forward equation, using the steady state as the initial condition g^0.

5. Given the evolution of the density of households, calculate the aggregate household supply of labour and demand for liquid and illiquid assets over time.

6. Update the path for wages, interest rates and house prices such that

(a) $w'_t = w_t - \xi_w \left(N_t + N^c_t - \int zl\,(a,b,z,t)\,d\mu_t \right)$, such that the wage falls (rises) when there is excess labour supply (demand).

(b) $\left(r^b_t \right)' = r^b_t - \xi_r \left(\int b\,d\mu_t - p^A_t C^{rent}_t + B^{gov}_t \right)$, such that the interest rate falls (rises) when there is an excess supply (demand) of liquid assets.

(c) $(p^a_t)' = p^a_t - \xi_a \left(IA_t - \delta A_t + \int \dot{a}_t + \dot{c}^{rent}_t\,d\mu_t \right)$, such that the price of houses falls (rises) when there is an excess supply (demand) of housing.

7. Repeat steps 2 through 6 until the sum of the changes in the three time path of prices $\left(w_t, r^b_t, p^a_t \right)$ converges.

B   Marginal Propensities to Consume

In this continuous time model I follow Achdou et al. (2017) in their definition of the marginal propensity to consume.

Definition 1. The marginal propensity to consume over a period τ for an individual with the state vector (a, b, z) is given by

$$MPC^x_\tau\,(a,b,z) = \frac{C_\tau\,(a,b+x,z) - C_\tau\,(a,b,z)}{x}, \tag{39}$$

where $C_\tau\,(a,b,z)$ is the sum of expected consumption of an individual household over the next τ periods.

$$C_\tau\,(a,b,z) = E\left[\int_0^\tau c\,(a_t,b_t,z_t)\,dt | a_0 = a, b_0 = b, z_0 = z \right] \tag{40}$$

This conditional expectation can be calculated using the Feynman-Kac formula. This formula links the conditional expectations of a stochastic process and the solution to partial differential equations.

C Alternate Model Specifications

C.1 Partial Borrowing Constraints

In the baseline version of the model I assumed that mortgage debt taken on by households is taken on a rolling basis. This implies that fluctuations in the price of housing will change the endogenous borrowing constraint faced by households. However, as discussed by Iacoviello and Pavan (2013), the assumption that the collateral constraint binds at all times is at odds with the reality that loan-to-value constraints only bind for agents that refinance. In practice, falls in the price of housing do not tighten the credit constraints of the majority who choose not to refinance. To better match the lagged correlation between measures of aggregate debt and house prices I relax this assumption to see how it affects the baseline results.

Concretely, I assume that the endogenous credit constraint applies directly to households that choose to adjust their stock of housing. For households who do not choose to adjust their stock of housing, I assume that stock of debt can be maintained at its current level even if the value of their housing stock falls due to a decline in house prices. This introduces an asymmetry into the relationship between house prices and the debt constraint. If house prices fall then households are not forced to immediately reduce their mortgage, but rather can maintain the size of the debt making repayments that cover the interest only. If house prices rise then households have the option of refinancing their mortgage and increasing their stock of debt. I assume that a household can costlessly refinance its current stock of housing. Thus the borrowing constraint becomes

$$b_t \geq -Ma_t p_t^a \qquad\qquad \text{if } I_{\alpha_t} = 1$$

$$(41)$$

$$b_t \geq \min\left\{-Ma_t p_t^a, b_t\right\} \qquad \text{if } I_{\alpha_t} = 0.$$

This changes both the steady state of the model, and its response to shocks around the equilibrium. While the price of housing is constant in the steady state, this alternative borrowing constraint loosens restrictions on households for whom the depreciation of their stock of housing would otherwise cause the debt constraint to bind. In equilibrium this relaxation allows households

to take on greater levels of debt due to the reduced probability of being credit constrained. Accordingly, the equilibrium interest rate is marginally lower.

The impact of this assumption on the estimated impact of the tax credit is marginal as house prices generally increase in response to the policy.

C.2 Heterogeneous Labour Supply

Previous research has indicated that households vary labour supply in response to changes in their balance sheet (Attanasio et al., 2012). In the baseline model, labour supply is invariant to a households' asset portfolio. Here I relax this assumption by considering a utility function in which the disutility from supplying labour is separable from consumption

$$u_t = \frac{\left(c_t^{1-\zeta} h_t^{\zeta}\right)^{1-\gamma}}{1-\gamma} - \phi z_t \frac{l_t^{1+\frac{1}{\eta}}}{1+\frac{1}{\eta}}. \tag{42}$$

Under this alternative function for flow utility, a household's optimal supply of labour will set by

$$z_t w_t \frac{\partial u_t}{\partial c_t} = -\frac{\partial u_t}{\partial l_t}$$

which in turn implies that each household's optimal labour supply $l_t = l_t\left(a, b, z\right)$ will vary according to the household's asset portfolio, though not with the productivity shock. To reflect the change in household preferences I re-calibrate the parameters within the utility function to match their original targets.

When I assume that labour supply varies with a households' assets, I create a second margin that households can adjust to respond to the change in debt that occurs when they adjust their stock of housing. Specifically, households increase their supply of labour in anticipation of, and after, buying additional stock. The reasoning behind this choice is the same as that of non-durable consumption, namely that households increase their labour supply to reduce the probability they will become bound by their credit constraint and reduce the time spent paying the higher interest rate on their mortgage. The same logic applies in reverse to households who anticipate decreasing their stock of housing.

Despite this second margin which households can use to offset fluctuations in the level of their non-durable consumption, the effect on marginal propensities to consume is quantitatively

49

small given reasonable calibrations for the Frisch elasticity of supply.

D   Rational Expectations and Instrumental Variables

Proposition 2. *Given a lagged instrument, an IV regression will produce a consistent estimator of the marginal effect of the expected change in the endogenous variable, ignoring the marginal effect of contemporaneous changes.*

Proof. Consider an endogenous variable which can be decomposed into a expected component based on information available at time $t - 2$ and a unexpected component, ξ_t

$$X_t = E_{t-2}X_t + \xi_t$$

Such that

$$E_{t-2}\left(\xi_t\right) = 0$$

Assume that in the true data generating process each component may have a different impact on the dependent variable, Y_t, such that

$$Y_t = \beta_1 E_{t-2}X_t + \beta_2 \xi_t + u_t$$

Now assume I have a instrument available, Z_{t-2}, that I wish to use to identify the marginal effect of the expected component, β_1, that is

Informative

$$E\left(Z_{t-2}E_{t-2}X_t\right) \neq 0$$

Valid

$$E\left(Z_{t-2}u_t\right) = 0$$

Only comprises information known at time $t - 2$ such that

$$E\left(Z_{t-2}\xi_t\right) = 0$$

Consider the regression of Y_t on X_t where X_t is instrumented with Z_{t-2}

$$\hat{\beta}_{IV} = \frac{\frac{1}{N} \sum_{i=1}^{N} (Z_{t-2})(Y_t)}{\frac{1}{N} \sum_{i=1}^{N} (Z_{t-2})(X_t)}$$

Therefore

$$\underset{N \to \infty}{\text{p} \lim} \hat{\beta}_{IV} = \frac{\frac{1}{N} \sum_{i=1}^{N} (Z_{t-2})(Y_t)}{\frac{1}{N} \sum_{i=1}^{N} (Z_{t-2})(X_t)}$$

$$= \frac{\underset{N \to \infty}{\text{p} \lim} \beta_1 \frac{1}{N} \sum_{i=1}^{N} (Z_{t-2}) E_{t-2}X + \beta_2 \frac{1}{N} \sum_{i=1}^{N} (Z_{t-2}) \xi_t + \frac{1}{N} \sum_{i=1}^{N} (Z_{t-2}) u_t}{\underset{N \to \infty}{\text{p} \lim} \frac{1}{N} \sum_{i=1}^{N} (Z_{t-2}) (E_{t-2}X_t) + \underset{N \to \infty}{\text{p} \lim} \frac{1}{N} \sum_{i=1}^{N} (Z_{t-2}) (\xi_t)}$$

By the law of large numbers

$$\underset{N \to \infty}{\text{p} \lim} \frac{1}{N} \sum_{i=1}^{N} (Z_{t-2}) E_{t-2}X = E((Z_{t-2}) E_{t-2}X) \neq 0$$

$$\underset{N \to \infty}{\text{p} \lim} \frac{1}{N} \sum_{i=1}^{N} (Z_{t-2}) (\xi_t) = E((Z_{t-2})(\xi_t)) = 0$$

$$\underset{N \to \infty}{\text{p} \lim} \frac{1}{N} \sum_{i=1}^{N} (Z_{t-2}) u_t = E((Z_{t-2}) u_t) = 0$$

Therefore

$$\underset{N \to \infty}{\text{p} \lim} \hat{\beta}_{IV} = \beta_1$$

Thus assuming that the instrument is valid and informative the IV estimator will be a consistent estimator of the marginal effect of the expected change in the independent variable. $\square$

E   Robustness of Empirical Analysis

In Table 10, I consider several alternative specifications for the regression of consumption on the change in household expectations. In Table 11, I show my results are robust over several alternative specifications for the regression of consumption on the change in household expectations. Finally, in Table **??**, I consider several alternative specifications for the regression of consumption

on the change in household expectations.

Table 10: Alternate regressions

	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	Baseline	Exclude food stamps	Pre-tax earnings	Continuously married only	Male heads only	Only mortgage debt	Include state unemployment
ΔDebt-Asset$_{i,t}$	-0.05 *	-0.07 **	-0.04	-0.08 *	-0.07 *	-0.11 **	-0.03 *
ΔFamily Size$_{i,t}$	0.01	0.02	0.01	0.01	0.01	0.01	0.01
r_{t-2}	-0.09***	-0.11***	-0.07*	-0.08***	-0.13***	-0.10**	-0.04*
$\Delta\log y_{i,t}$	-0.01	0.02	-0.01	-0.02	-0.03	-0.00	0.01
$\Delta\log W_{i,t}$	-0.00	0.01	-0.00	-0.00	0.01	0.02	-0.01
ΔState Unemployment$_{i,t}$							-0.05
Observations	37,560	37,560	37,560	21,242	26,185	37,560	37,560

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 11: Changes in household expectations

		Baseline	Exclude food stamps	Pre-tax earnings	Continuously married only
	No expectation to Definitely	-0.073^{***}	-0.088^{***}	-0.069^{***}	-0.053^{*}
Change in household expectations	Uncertain to Definitely	-0.063^{***}	-0.060^{**}	-0.055^{***}	0.013
	Probably to Definitely	-0.048^{***}	-0.052^{**}	-0.072^{***}	-0.022
	Definitely to Definitely	-0.057^{***}	-0.046^{*}	-0.007	-0.07

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

E.1 Durable good consumption

In Table 12, I consider the inclusion of durable good consumption. Similar to Benmelech et al. (2017) I find that consumption of durable goods, in particular house furnishings, increases when a household purchases a house. However I find that total consumption still falls when durable goods are included in the consumption regression. This difference may be due to omitted variable bias such as contemporaneous shocks to financial wealth or income which are controlled for under the instrumental variable method. I also consider the possibility that the future expectations regarding the decision to purchase a house are correlated with changes in household size. To test if this correlation impacts the results I re-run the regressions on sub-samples of the population

limited to those who do not report a change in family size or households over the age of 40 for whom the average fertility rate is significantly lower. Table 13 shows that these alternative samples do not materially changes the main results of the paper.

Table 12: Alternate regressions - durable goods

	All Buyers			First-time buyers			Repeat Buyers		
	Non-durable	Durable	Total	Non-durable	Durable	Total	Non-durable	Durable	Total
$\text{Buy}_{i,t}$	-0.11 ***	0.48 ***	-0.08 ***	-0.07 *	1.922 ***	-0.01	-0.62 *	-0.02	-0.02
$\Delta\text{Family Size}_{i,t}$	0.02	0.01	0.01	0.02	0.01	0.01	0.01	0.01	0.01
r_{t-2}	-0.07***	-0.01***	-0.08*	-0.08***	-0.03***	-0.10**	-0.04*	-0.01***	-0.08*
Observations	37,560	31,560	31,560	37,560	31,560	31,560	37,560	31,560	31,560

$^{*}\ p < 0.05,\ ^{**}\ p < 0.01,\ ^{***}\ p < 0.001$

Table 13: Alternate regressions - household size

	All Households		Age > 40		Family size unchanged	
	Non-durable	Total	Non-durable	Total	Non-durable	Total
$\text{Buy}_{i,t}$	-0.11 ***	-0.08 ***	-0.15 ***	-0.14 ***	-0.12 ***	-0.10 ***
$\Delta\text{Family Size}_{i,t}$	0.02	0.01	0.01	0.01	0.01	0.00
r_{t-2}	-0.07***	-0.08*	-0.08***	-0.10**	-0.04*	-0.08*
Observations	37,560	31,560	27,668	17,261	33,168	28,810

$^{*}\ p < 0.05,\ ^{**}\ p < 0.01,\ ^{***}\ p < 0.001$