

Consumption Response to Aggregate Shocks and the Role of Leverage

Agnes Kovacs, May Rostom, Philip Bunn *

July 11, 2018

Abstract

This paper investigates the relationship between mortgage leverage and consumption around the 2008 financial crisis. Using data from the UK's Family Expenditure Survey and Wealth and Asset Survey, we first show that high-leveraged households made larger cuts to consumption following the financial crisis, and this was largely driven by young households. Second, using a life-cycle framework, we investigate the channels by which high-leveraged households may have reduced consumption by more than others. Our key finding is that credit supply tightening is the main driver of the empirical co-movement between pre-crisis leverage and consumption growth after 2008.

Keywords: life-cycle models, consumption, household leverage, debt, financial crisis

JEL classification: D10; D11; D14; E21

*Kovacs: University of Oxford; Rostom: University College London and Bank of England; Bunn: Bank of England. We thank Orazio Attanasio, Richard Blundell, Thomas Crossley, Antonio Guarino, Patrick Moran, Lars Nesheim, Cormac O'Dea, and Morten Ravn for helpful comments and advice. We also thank seminar participants at RES (Bristol) conference, EEA (Lisbon) conference, IAAE (Montreal) conference, UCL PhD Summer conference, Bank of Canada, Bank of England, Bank for International Settlements, and University of Ottawa for their feedback. Any views expressed are solely those of the author(s) and so cannot be taken to represent those of the Bank of England or to state Bank of England policy. This paper does not necessarily represent the views of the Bank of England or members of the Monetary Policy Committee or Financial Policy Committee. All errors are our own.

1 Introduction

UK Household consumption fell sharply during the 2008 financial crisis, by 4 percent overall. However, the fall was even more drastic for households with high levels of mortgage: over 10 percent on average. To understand the effects of the crisis it is essential to explain this first-order response. In this paper, we focus on understanding and quantifying the different channels through which mortgage status affected households' consumption behaviors in the last crisis.

Using household-level data from the UK's Family Expenditure Survey (FES) and from the Wealth and Asset Survey (WAS), we first document the empirical relationship that households with a mortgage cut their consumption by more than households with no mortgage. We show that the higher a household's pre-crisis leverage (as measured by their loan-to-income ratio) the more they cut their consumption in the crisis period. Moreover, we find that the relationship between leverage and consumption growth varies significantly with the observed time period and with age. On the one hand, this co-movement seems to be present in the crisis period only. When we examine non-recessionary periods, we find no significant impact of household leverage on consumption behavior. On the other hand, young households cut their consumption by roughly twice as much as old households for the same level of leverage over the crisis.

Having documented these empirical irregularities, we then turn to understanding the drivers that led indebted households to cut their consumption by more than other households during the crisis. We consider four channels by which the recession may have directly affected households: a fall in the level of income, rising uncertainty around future income, a reduction in the supply of credit, and a decline in house prices. As we argue below, all four channels can generate a positive co-movement between leverage and consumption growth.

Lack of data makes it impossible to test the channels empirically, therefore we build a life-cycle model with a realistic mortgage market to assess the quantitative importance of these four mechanisms. Our simulation results show that the most important driver of the empirical co-movement between pre-crisis leverage and consumption growth in the data was the exogenous reduction in the supply of credit after 2008.

Consider how the four channels affect households with different pre-crisis leverage ratios.

The first channel works through a negative shock to income. When households face a fall in their income level, they cut their consumption. However, the size of the cutback may differ across households not just because of where they are in their life-cycle but also

because of liquidity constraints. After a negative permanent income shock of the same size, younger households may cut their consumption by more than older households as they are likely to be more liquidity constrained. As young households are typically more leveraged than old households, we may also observe that more leveraged households decrease their consumption by more than others after a negative income shock.

The second channel plays a role when the recession is associated with greater uncertainty around future income. As uncertainty increases, households increase their precautionary savings, and hence consume less. Again, depending on where the households are in their life-cycle and consequently what their leverage is, they could quantitatively behave very differently.

The third channel works through an exogenous reduction in the supply of credit. Since the crisis, banks significantly decreased the acceptable mortgage loan-to-value ratios (LTV) at which they were willing to lend, which made it increasingly difficult for households to finance consumption by topping up mortgages. As a result, households close to the acceptable leverage threshold were unable to withdraw additional equity from their housing wealth. This change may affect high-leveraged households the most, leading their consumption to fall by much more than other households.¹

The final channel we consider is a fall in house prices, which is very similar to the credit supply channel. As house prices fall, mortgagors face a reduction in their housing wealth, which increases their LTV ratios. These high LTV households would then be unable to withdraw additional equity from their house, which consequently may lead them to decrease their consumption by more than other households.

Although the four channels have a similar qualitative effect in encouraging a positive co-movement between leverage and consumption growth, building a structural model allows us to investigate these channels independently. In particular, we use model simulations to show which channel leads to a similar consumption response to what we observe in the data. We find that the observed pattern of consumption in the recent recession is best explained by a model in which the recession is associated with a credit supply shock.

The literature linking consumption growth to changes in incomes and house prices is extensive. For example, Attanasio and Weber (1994) and Attanasio et al. (2012) show that consumption growth for young and old households varies by the type of shock. Positive income (and income expectations) shocks drive consumption growth among young households, whereas positive house price shocks are the main driver of consumption growth for old households.

¹It would also affect first time buyers. Whereas before the crisis first time buyers only needed to save 5% of the value of a house, the required down payment increased significantly after the crisis hit.

There is also a substantial literature addressing the question on whether housing wealth is used as buffer saving. Skinner (1996) shows that younger households adjust their consumption when there is a change in their housing wealth. Hurst and Stafford (2004) show that households use their housing wealth to refinance their mortgage in order to smooth consumption over their life-cycle. Campbell and Cocco (2007) find a correlation between house prices and predictable changes in consumption, which is more pronounced among credit constrained households. The study by Berger et al. (2016) also confirms that the size of housing price effects on consumption may be substantial and sensitive to the level of household debt.

There are, however, fewer studies that analyze the relationship between mortgage and consumption behavior. Dynan (2012) uses household level data to show that households with high LTV have experienced larger declines in spending during the crisis, after controlling for wealth effects and other factors. Mian, Rao, and Sufi (2013) show that the decline in US consumption following the crisis was greater in counties with higher leverage prior to the crisis. Baker (2013) finds that spending by highly leveraged households has been more sensitive to income fluctuations and that it is likely to have increased the depth of the recent recession.

More recently, Cloyne and Surico (2017) use UK data to show that households with a mortgage adjust their consumption in response to tax changes by much more than households with no mortgage. The authors argue that the lack of net liquid wealth and the transaction costs associated with liquidating housing assets make indebted households more liquidity-constrained. An income shock, therefore, may elicit a bigger response to consumption for mortgagors.² Kaplan, Mitman, and Violante (2017) study the role of housing markets and mortgages. They find that access to credit affects home-ownership; and that changes in housing wealth (through changes in house prices) account for about half of the changes in non-durable consumption. Finally, the work by Wong (2015) shows that different types of households may respond differently to shocks. In particular, she finds that the consumption of young, leveraged households is more responsive to monetary policy shocks than other types of households.

The rest of the paper is organized as follows. In Section 2, we describe the data and motivation for our study by presenting the empirical relationship between leverage and

²While this seems supportive of our income shock channel, there is an important difference between their analysis and ours. Cloyne and Surico (2017) examine the effects of tax changes on income, but the majority of their observed tax changes were in fact designed to boost income. We are specifically interested in the impact of a negative income shock in the period after 2008, whereas their analysis is for the period before the crisis. This difference is crucial as a household's response to positive and negative income shocks is likely to be asymmetric given that the first relaxes, while the second triggers a liquidity constraint.

consumption growth. In Section 3, we give more elaborate evidence of this relationship using a regression analysis. In Section 4, we detail our theoretical life-cycle model together with our solution strategy and model calibration. In Section 5, we present the results. In Section 6, we discuss the implications of our analysis and conclude the paper.

2 Consumption and Leverage

Aggregate data from the UK statistical agency, the Office for National Statistics (ONS), show that household consumption was growing at an annual rate of 4% per year in the decade running up to the crisis. It collapsed in 2009 by a massive 7 percentage points—from a peak of roughly 3% a year to -4% in the depth of the crisis. Since 2009 (and at least until 2012), it grew at rates below its historical average (Figure 1).³




Figure 1: Annual Real Consumption Growth

Source: Office for National Statistics. Final household consumption expenditure, chained-volume measures, excluding imputed rents (housing).

These aggregate data potentially mask a large degree of heterogeneity across households for a number of reasons— for example because of differences in financial position, age or education. In Figure 2 we disaggregate consumption data and plot annual consumption growth for households with different housing tenures. The purple line plots renters; the green line plots homeowners with no mortgage debt; the orange line plots homeowners who have a mortgage loan of less than twice their income; while the blue line plots

³The level of consumption was also very weak and has not recovered since.

homeowners who have a mortgage loan of more than twice their income. As Figure 2 shows, households with different housing tenures differed in their consumption behaviors in the two decades preceding 2012. The most striking difference however, arose over the great recession. While all households reduced their consumption in the recession, the biggest—and so far unprecedented—consumption adjustment came from those with the highest leverage. Compared to the 4% average drop in consumption in 2009, households with a mortgage loan (or debt) to income ratio (LTI) of two or more decreased their consumption by over 10%.

Our goal in this paper is to understand the mechanisms that caused the behavior we see in Figure 2. Simply put, what led high-leveraged households to cut their consumption by much more than low-leveraged ones? In this, and the following section, we focus on a deeper empirical analysis of the available data.




Figure 2: Annual Real Consumption Growth by Tenure

Source: Family expenditure survey (FES). Consumption is total weekly real non-housing consumption per household. Low LTI households are those with a mortgage loan-to-income ratio of less than 2; high LTI households are those with ratios of 2 or more. Homeowners with no mortgage account for 32% of households; Low LTI and High LTI both account for 19% each; and Renters account for the remaining 30% share. FES data are scaled to match the ONS National Accounts.

2.1 Data

We use data from the Family Expenditure Survey (FES), a micro data-set containing detailed information on household expenditures of 5,000 to 6,000 households per year. A two-week expenditure diary accompanies the questionnaire. The FES provides the best

household-level consumption data in the United Kingdom. It also contains information on outstanding levels of mortgage debt from 1992. However, our main analysis predominantly covers the period between 2006 to 2012. We focus on households of working age, where the head is aged between 21 and 69. Waves of the FES also switch between calendar and fiscal year collection multiple times, so to maintain consistency, we convert all data to calendar years.

We use weekly non-housing expenditure as our measure of consumption⁴, and deflate all data using the National Accounts consumption deflator. We focus only on secured debt as the FES has only scant data on unsecured debt. As mortgage debt accounts for the majority (>80%) of all debt, excluding unsecured debt is unlikely to change the big picture message.

The Survey contains no information on assets. To address this, we merge the FES with data from the Wealth and Assets Survey (WAS) at the cohort level. The WAS is a panel survey covering 20-30,000 households in each wave. The survey began in 2006 and we use its first three waves for our analysis: Wave 1 from mid- 2006 to mid-2008, Wave 2 from mid-2008 to mid-2010 and Wave 3 from mid-2010 to mid-2012.⁵ Specifically, we merge mean values of real housing wealth and real gross non-housing financial wealth⁶. This allows us to construct measures of LTV, in line with how most of the literature defines leverage. While we use this definition throughout the paper to support our results, the lack of asset data availability before 2006 mean that we cannot rely on it for regressions we run prior to 2006 (for robustness checks). As such, our main definition of leverage is LTI. This is also an appropriate measure to use for the UK for two reasons. First, this is a ratio that UK banks rely on as a measure to ration access to credit (although it is secondary to LTV). Second, it allows us to test the impact of falls in income on household leverage. Nevertheless, we find similar effects at the cohort level regardless of which method of leverage we use.

Unfortunately, because the FES are an annual time-series of cross-sections, we cannot control for individual-level heterogeneity or track changes over time. Instead, we construct a synthetic panel, using date of birth, as first suggested in Deaton (1985a) and in line with what others using the FES have done in the past (see Attanasio and Weber

⁴The methodology for calculating housing consumption in the FES is not consistent with that used in the National Accounts. For homeowners, the FES only measures mortgage payments rather than using a measure of imputed rents like that of the National Accounts. In common with papers who use this micro-data, we scale incomes, consumption and debt by the ratio of its FES total for all households to its National Accounts aggregate.

⁵While the time periods from the WAS wave do not exactly match the calendar time periods in our analysis, they nonetheless provide a reasonable approximation as assets tend to change slowly over time. Using data from part of a WAS wave risks the data not being fully representative.

⁶This excludes bank deposits.

(1994), Attanasio, Banks, and Tanner (2002), or Campbell and Cocco (2007) as examples). The premise of this approach is that while we cannot track individual households over time, we can track groups of households that share similar characteristics. Typically, households are grouped in cohorts using their year of birth (as age is exogenous). However, depending on the nature of the question being asked, other groupings have also been constructed, such as years of education or housing tenure status. This method allows us to condition debt at its pre-crisis levels and to consider the effects on different households, albeit at a lower level of disaggregation than is possible at the household level.

We use two different cohort definitions to ensure the reliability of our results. The first definition uses year of birth to construct the cohorts. For example, everybody born in 1956 belongs to one cohort. This is the cleanest definition of a cohort because birth year is deterministic—people cannot move across cohorts over time. However, it also pools people with and without mortgage debt making it harder to tease out the relationship between debt and spending patterns. To discriminate between these, we consider a second definition: splitting each birth year by whether or not they have a mortgage⁷. The main disadvantage of the second cohort definitions is that the sample split is endogenous (households choose their mortgage status). In Appendix A.4 we use the method suggested by Moffitt (1993) to construct predicted values of leverage. Our results continue to hold.

We pool two years together to construct our time periods. We define 2006/2007 as the pre-crisis period and 2009/2010 as the period thereafter (we later extend the analysis to 2011/2012 too)⁸. This pooling makes the time periods of the data more comparable to the WAS waves, where each wave also spans two years. It also allows us to increase the number of data points that is used to create each cohort level observation.

2.2 Descriptive Statistics

When we think about the relationship between leverage and consumption growth, we need to focus on three variables that determine this connection: income, mortgage and consumption. These variables vary both over the life-cycle and across time. In this

⁷There is a tradeoff between the mean number of observations used to create each cohort cell and the total number of observations overall. We drop cohorts with less than 50 observations per cell.

⁸2009-2012 was a period of prolonged weakness: Aggregate levels of consumption, the total wage bill, and house prices all fell to their lowest level since a decade ago; unemployment rose to its highest level since the mid-1990s; and high LTV products became very difficult to access. Indeed, very high LTV products were taken completely off the market and were impossible to access during this time period.

section, we examine how they evolved between 1992 and 2012 using cohort analysis.

Figure 3 plots these three key variables between 1992-2012. Panel (a) plots mean real incomes, panel (b) plots mean real consumption and panel (c) plots mean outstanding mortgage debt. Cohorts are defined by 10-year-of-birth intervals. Each line shows the mean for a variable of interest for a given cohort. For example, the purple line in panel (a) represents the mean income for households that were born between 1940 and 1949. The left-most data point of the line is from the earliest available survey in 1992, whereas the right-most data point of the line is from the latest wave in 2012⁹. Moving from left to right shows the evolution of income, consumption and debt for different generations as they age over time.

All of the plotted variables - income, consumption and outstanding mortgage debt - are hump-shaped over the life-cycle and peak before retirement, consistent with the theory that households accumulate mortgage debt when young and convert it to assets as they age, thereby smoothing consumption.


Panel (a) shows how the incomes of all cohorts have flattened out, falling from about 2007 onward. In fact, some younger cohorts earned less than their immediate preceding cohort when they were at the same point in their life-cycle. This pattern is also mirrored in Panel (b). However, compared to income growth, consumption growth has been less rapid both between and across cohorts. Turning to mortgage debt, Panel (c) shows that younger cohorts have much more debt compared to older cohorts when they were of the same age. Furthermore, debt grew at much faster rates than income. The financial liberalization of the 1990's, which made credit more readily available, may have played a significant role in the increase of mortgage debt among young households.

Since the rise in mortgage debt may also be a reflection of the rapid rise in house prices over the last two decades, in Figure 4 we plot real outstanding mortgage debt and real housing wealth. We use repeated cross-section data from three different periods: 1995, 2007 and 2012¹⁰. These graphs suggest that the beneficiaries of the rising house prices were the older cohorts. Even though younger cohorts took on the lion's share of the increase in mortgage debt, older cohorts had the biggest increase in housing wealth. One might therefore expect big differences in household spending growth across households of different age groups.

⁹The youngest and oldest cohorts have shorter lines on the charts as they were either younger than 21 or older than 70 in the earlier/later waves of the survey.

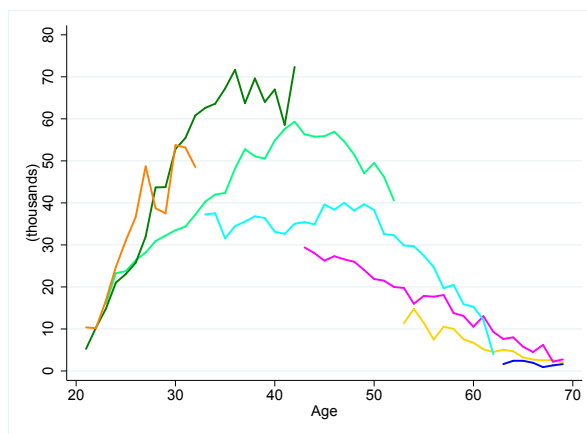
¹⁰As no wealth information exists in the FES, data for 2007 and 2012 lines in Panel(b) are taken from the 2006/2007 and 2011/2012 WAS waves. As WAS only dates back to 2006/2007, data for 1995 are taken from Bunn and Rostom (2014) using the British Household Panel Study (BHPS). To ensure data accuracy of these spliced datasets, we plotted the 2006/2007 WAS line against the 2005 data taken from the BHPS. We get the similar results.

We report summary statistics from the FES and WAS in the Appendix.



(a) Real weekly household income


(b) Real weekly household consumption



(c) Real mortgage debt

Figure 3: Life-cycle behaviors of key variables by age across cohorts

Source: Family Expenditure Survey (FES) from 1992 - 2012, for head of household aged 21-69 at the time of the survey. Each line represents a 10-year date of birth cohort. Incomes are net of tax. Consumption excludes housing expenditures. Mortgage debt is outstanding mortgage debt. Data are deflated using the 2011 consumption deflator.



(a) Real outstanding mortgage debt

(b) Real housing wealth

Figure 4: Average household debt and housing wealth by age across years

Source: Family Expenditure Survey (FES) for Panel (a). British Household Panel Study (BHPS) and Wealth & Assets Survey (WAS) for Panel (b). In Panel (b), data for 1995 are taken from Bunn and Rostom (2014); whereas data for 2007 and 2012 are from WAS Waves 1 and 3 respectively. All data are for head of household aged 21-69 at the time of the survey. Data are deflated using the 2011 consumption deflator.

3 Empirical Determinants of the Link Between Consumption and Leverage

After the graphical description of the relationship between leverage and consumption growth in the previous section, in this section, we give a more elaborate analysis of this observation. We control for different factors which we think might affect the observed connection between leverage and consumption growth, but these variables will undoubtedly still be endogenous. Therefore, the goal of this section is to establish the existence of this co-movement rather than attribute causality. Our regression strategy is motivated by the availability of data: we run our cohort regressions on the change between 2006/2007 and 2009/2010. In Appendix A.5 we extend our analysis by substituting 2009/2010 with 2011/2012 instead. Our results continue to hold.

3.1 Baseline Regression

In our baseline estimation, we are interested in exploring Figure 2 further. In particular, whether we can still see a relationship between leverage and consumption in a simple, reduced form estimation, without aiming to identify these coefficients. Instead, we want

to observe whether leverage, as measured by the LTI, before the crisis (2006/2007) was associated with a change in consumption during the period of weak growth (2009/2010). We estimate the following cross-sectional equation at the cohort level:

$$\Delta c_{it} = \alpha_0 + \alpha_1 LTI_{it-1} + \alpha_2 \Delta y_{it} + \alpha_3 \Delta w_{it} + \alpha_4 x_{it} + e_{it} \quad (1)$$

Where Δc_{it} is the change in log non-housing consumption for cohort i in 2009/2010, denoted by t . LTI_{it-1} is the ratio of outstanding mortgage debt to after-tax income in 2006/2007, denoted by $(t - 1)$. y_{it} is log after-tax labor income; w_{it} is a vector of log housing wealth and log financial wealth; while x_{it} is a vector of demographic characteristics, such as change in household composition. As discussed before, FES contains no information on wealth, hence wealth data is taken from the Wealth and Asset Survey (WAS). Notice that equation (1) is similar to the one used by Dynan (2012).

We are interested in the coefficient on the pre-crisis leverage, α_1 , since it shows the effect of leverage on consumption growth after controlling for income and wealth. Considering Figure 2 again, we expect this coefficient to be negative: higher pre-crisis leverage is expected to lead to bigger falls in consumption growth.

Table 1 reports the regression results from equation (1). Each column represents one of the cohort definitions as specified in Section 2.1. In both columns, the coefficients on pre-crisis LTI are negative and statistically significant. They are also broadly similar in magnitude across all definitions. On average, we find that a one unit increase in the pre-crisis leverage (e.g. LTI of 3 instead of 2) is associated with about a 2.6 percentage-point decrease in consumption growth.

Notice that ideally, we would also like to control for variations in permanent income, income uncertainty, down payment requirement and house prices when running equation (1). Interacting these variables with pre-crisis LTI would help us get closer to identifying the driver of the observed co-movement between leverage and consumption growth. Unfortunately because of data availability (e.g. we have no consistent household panel data for this time period), we cannot check this using the micro-data. (In the following Section we discuss how we use the structural model allows for a more detailed counterfactual analysis).

| Cohort definition | Year of birth | Mortgagor / Non-Mortgagor |
|------------------------------|----------------------------|------------------------------|
| | [1] | [2] |
| Dependent var: ΔC_t | | |
| LTI_{t-1} | -0.027** (0.014) | -0.026*** (0.008) |
| ΔY_t | 0.653*** (0.132) | 0.546*** (0.117) |
| $\Delta Housing\ wealth_t$ | -0.022 (0.068) | 0.073** (0.035) |
| $\Delta Financial\ wealth_t$ | 0.002 (0.022) | 0.004 (0.023) |
| Constant | -0.025 (0.023) | -0.028** (0.012) |
| Controls | Yes | Yes |
| Observations | 46 | 77 |

Robust standard errors in parentheses, *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. ΔC_t is the change in log non-housing consumption in 2009/2010, denoted by t . ΔY_t is the change in after-tax income, net of mortgage interest payments at t . LTI_{t-1} is the ratio of mortgage debt to annual income in 2006/2007, denoted by $t - 1$. $\Delta Housing\ wealth_t$ and $\Delta Financial\ wealth_t$ are changes in housing wealth and total financial wealth respectively. Controls are the change in the number of adults and children. Each column represents a synthetic panel regression from 2 pooled years, constructed using either [1] year of birth, or [2] year of birth, split by mortgage status.

Table 1: Loan-to-Income

3.2 Regressions by Age

Household debt accumulation is disproportionately skewed towards younger households. Also, narrative evidence suggests that the Great Recession was harsher on younger households than it was on older households, as they had bigger falls in their income

and were more likely to become unemployed. Tighter credit conditions may have also affected younger households more than older households.

| Sample split: | Young | Old |
|------------------------------|-----------------------------|---------------------------|
| | [1] | [2] |
| Dependent var: ΔC_t | | |
| LTI_{t-1} | -0.035*** (0.010) | -0.020* (0.010) |
| ΔY_t | 0.327* (0.157) | 0.756*** (0.141) |
| $\Delta Housing\ wealth_t$ | 0.082* (0.039) | -0.022 (0.079) |
| $\Delta Financial\ wealth_t$ | 0.016 (0.037) | -0.001 (0.030) |
| Constant | -0.007 (0.024) | -0.041** (0.018) |
| Controls | Yes | Yes |
| Observations | 26 | 51 |

Robust standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1. Sample is split by whether the household head is younger than 40 years or otherwise. ΔC_t is the change in log non-housing consumption in 2009/2010, denoted by t . ΔY_t is the change in after-tax income, net of mortgage interest payments at t . LTI_{t-1} is the ratio of mortgage debt to annual income in 2006/2007, denoted by $t - 1$. $\Delta Housing\ wealth_t$ and $\Delta Financial\ wealth_t$ are changes in housing wealth and total financial wealth respectively. Controls are the change in the number of adults and children. Regressions are from a synthetic panel using 2 pooled years, where the panel is constructed using year of birth and mortgage status.

Table 2: Results by Age

In order to test whether the changed conditions during the Great Recession had a bigger effect on young household's consumption behaviors than on the old, we split our sample into two by age. We define young households as those who are younger than 40 years of age and old households as those 40 years or more. Our results are reported in Table 2. The coefficients on LTI_{t-1} are 1.75 times greater for young households than for old households, suggesting that young households with high leverage cut back consumption the most during the great recession.

3.3 Regression in Non-Crisis Periods

One challenge to our results is that the effect we are finding may not be specific to only recessions. To counter that, we repeat our analysis using three non-crisis time periods between $(t - k)$ and (t) . The first time period is between 1997/1998 and 2000/2001; the second time period is between 2000/2001 and 2003/2004; and the third is between 2003/2004 and 2006/2007. Table 3 reports results for these placebo tests¹¹. In all cases we find no effect— the coefficients are small and insignificant.

We interpret this to mean that the drivers that lead indebted households to change their consumption are specific to recessionary shocks. We discuss the effects of these drivers in the next section.

¹¹The equations do not include wealth variables as WAS did not exist before 2006.

| Time period, t : | 2000/2001 | 2003/2004 | 2006/2007 |
|-----------------------------|--------------------------|-------------------------|-------------------------|
| | [1] | [2] | [3] |
| Dependent var: ΔC_t | | | |
| LTI_{t-1} | -0.002 (0.008) | 0.010 (0.008) | 0.006 (0.008) |
| ΔY_t | 0.611*** (0.061) | 0.517*** (0.051) | 0.459*** (0.125) |
| Constant | 0.014 (0.014) | -0.013 (0.011) | -0.032** (0.014) |
| Controls | Yes | Yes | Yes |
| Observations | 79 | 79 | 79 |

Robust standard errors in parentheses, *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Each column represents a regression for a different time period. ΔC_t is the change in log non-housing consumption in t ; where t is [1] 2000/2001, [2] 2003/2004, or [3] 2006/2007 for each column respectively. ΔY_t is the change in after-tax income, net of mortgage interest payments at t . LTI_{t-1} is the ratio of mortgage debt to annual income in 2006/2007, denoted by $t - 1$; where $t - 1$ is [1] 1997/1998, [2] 2000/2001, [3] 2003/2004 for each column respectively. Controls are the change in the number of adults and children. No controls for wealth, as data does is not available before 2006. Regressions are from a synthetic panel where the panel is constructed using year of birth and mortgage status.

Table 3: Non-Crisis Period

3.4 Sensitivity analysis

In this section, we test different specifications to check the robustness of our results. We summarize those below, although all robustness checks are detailed in Appendix A.2.

First, we use a more commonly used definition of leverage measured by the mortgage loan-to-value ratio (LTV), in keeping with the way other studies commonly define leverage. This specification allows us to compare our results to other studies in the literature. Using the housing wealth data from WAS, we construct loan-to-value ratios and estimate a similar equation to equation (1) but substitute LTIs with LTVs:

$$\Delta c_{it} = \beta_0 + \beta_1 LTV_{it-1} + \beta_2 \Delta y_{it} + \beta_3 \Delta w_{it} + \beta_4 x_{it} + e_{it} \quad (2)$$

The results are reported in Table 4. On average, we find that a ten unit increase in the pre-crisis leverage, as measured by LTV, led to a 14-percentage-point decrease in consumption growth. When we split the sample by young and old, we also find that leverage had a much bigger effect on young household's consumption growth compared to old households. More specifically, we find that a ten unit increase in the pre-crisis LTV led to an 18-percentage-point decrease in consumption growth for young households (compared to only 12-percentage-points for old households). Hence, our conclusion is similar to the case where we measure leverage by loan-to-incomes: there is a negative connection between leverage and consumption growth even after controlling for income and wealth. Results from equation (2) using our year-of-birth cohort definition can be found in Appendix A.3.

Our second robustness check is related using mortgage status to create the cohorts. The main criticism of this definition is that mortgage status itself is a choice variable, so the construction of the panel does not just rely on deterministic variables.

Following Attanasio, Banks, and Tanner (2002), we use predicted mortgage status (instead of actual status) to construct an alternative panel specification¹². The detailed method and results are discussed in Appendix A.4. In summary, we get the same results using the predicted variables as we would with the actual variables. This result is not that surprising, since data from the WAS suggest that around 90% of the households who had a mortgage in the first wave of the survey, between mid-2006 and mid-2008, still had a mortgage two years later. We can therefore assume that over shorter periods - such as when looking at changes over the financial crisis period - mortgage level and status are fixed.

¹²We also use predicted debt for the LTI ratio.

| Sample split: | Everyone | Young | Old |
|------------------------------|-----------------------------|-----------------------------|---------------------------|
| | [1] | [2] | [3] |
| Dependent var: ΔC_t | | | |
| LTV_{t-1} | -0.142*** (0.039) | -0.177*** (0.049) | -0.118* (0.062) |
| ΔY_t | 0.551*** (0.117) | 0.328* (0.164) | 0.755*** (0.142) |
| $\Delta Housing\ wealth_t$ | 0.074** (0.035) | 0.083** (0.039) | -0.020 (0.079) |
| $\Delta Financial\ wealth_t$ | 0.005 (0.023) | 0.019 (0.036) | -0.001 (0.030) |
| Constant | -0.028** (0.012) | -0.007 (0.024) | -0.042** (0.018) |
| Controls | Yes | Yes | Yes |
| Observations | 77 | 26 | 51 |

Robust standard errors in parentheses, *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Sample is split by whether the household head is younger than 40 years or otherwise. ΔC_t is the change in log non-housing consumption in 2009/2010, denoted by t . ΔY_t is the change in after-tax income, net of mortgage interest payments at t . LTV_{t-1} is the ratio of mortgage debt to housing wealth in 2006/2007, denoted by $t - 1$. $\Delta Housing\ wealth_t$ and $\Delta Financial\ wealth_t$ are changes in housing wealth and total financial wealth respectively. Controls are the change in the number of adults and children. Regressions are from a synthetic panel using 2 pooled years, where the panel is constructed using year of birth and mortgage status.

Table 4: Loan-to-Value

| Sample split: | Everyone | Young | Old |
|------------------------------|----------------------------|-----------------------------|--------------------------|
| | [1] | [2] | [3] |
| Dependent var: ΔC_t | | | |
| $LTI_{t-1} > 2$ | -0.049** (0.020) | -0.093*** (0.031) | -0.010 (0.028) |
| ΔY_t | 0.529*** (0.118) | 0.310* (0.157) | 0.730*** (0.148) |
| $\Delta Housing\ wealth_t$ | 0.086** (0.035) | 0.086** (0.040) | 0.033 (0.076) |
| $\Delta Financial\ wealth_t$ | 0.008 (0.025) | 0.018 (0.041) | 0.004 (0.030) |
| Constant | -0.038*** (0.012) | -0.011 (0.024) | -0.048*** (0.017) |
| Controls | Yes | Yes | Yes |
| Observations | 77 | 26 | 51 |

Robust standard errors in parentheses, *** p<0.01, ** p<0.05, * p<0.1. In columns 2 and 3, the sample is split by whether the household head is younger than 40 years or otherwise. ΔC_t is the change in log non-housing consumption in 2009/2010, denoted by t . ΔY_t is the change in after-tax income, net of mortgage interest payments at t . $LTI_{t-1} > 2$ is an indicator variable for whether LTI was > 2 in 2006/2007, denoted by $t - 1$. $\Delta Housing\ wealth_t$ and $\Delta Financial\ wealth_t$ are changes in housing wealth and total financial wealth respectively. Controls are the change in the number of adults and children. Regressions are from a synthetic panel using 2 pooled years, where the panel is constructed using year of birth and mortgage status.

Table 5: LTI > 2

Third, we run an alternative estimation with an indicator variable for whether a cohort has an LTI greater than two¹³. This gets us closest to what we show in Figure 2. The results, which are reported in Table 5, show that, overall, consumption growth for cohorts who have an LTI greater than two is associated with a 4.9 percentage point fall

¹³The mean loan-to-income value at the cohort level is equal to two.

between 2006/2007 and 2009/2010. But the effect is very stark when households are split by age— young households decreased their consumption growth by 8 percentage points more than old households. In Appendix A.6, we take this analysis further by testing for non-linearities in leverage by creating additional flags for cohorts with LTIs greater than 2.5 and 3. We find that the greater the LTI threshold, the greater the consumption response.

Fourth, we extend the post-crisis time period to 2011/2012. We find that extending our estimation period gives us a similar size and significance of the coefficients on the loan-to-income ratio, suggesting the effects were persistent.

So far, we have used micro-data at the household and cohort level to show the relevance of leverage status on households' consumption behaviors during the financial crisis. We demonstrated three empirical facts of consumption growth. First, both low- and high-leveraged households cut their consumption in response to the crisis. Second, high-leveraged households cut their consumption by more than low-leveraged ones. Third, this was mostly driven by young households.

4 A Life-Cycle Model with Mortgages

In a regression analysis it is impossible to disentangle how the consumption behavior of high- and low-leveraged households is affected by potentially different shocks around the recession. We cannot convincingly test whether our empirical findings were mainly driven by income shocks, credit market shocks or house price shocks. For this reason, we build a theoretical model and simulate the behavior of a large number of households. The model enables us to disentangle the effects of different channels, and to show which of them can lead to consumption patterns similar to those we highlighted in the previous section, i.e. the three empirical facts of consumption growth.

4.1 Model structure

We start with a simple model of life-cycle consumption and savings in a dynamic stochastic framework. Households maximize their present discounted lifetime utility, which depends on nondurable consumption and the consumption of a flow of housing services. They can move resources between periods by investing in either a one-period bond or in housing, which also provides a flow of housing services. They are only allowed to have collateralized debt, and only housing can serve as collateral. Households face un-

certainty in two dimensions: idiosyncratic uncertainty over labor income and aggregate uncertainty over house prices.

Utility function. Households have a CRRA utility function in the composite good, which in turn is a Cobb-Douglas aggregate of nondurable consumption, C_t and housing services, S_t .

$$u(C_t, S_t) = \frac{(C_t^\alpha S_t^{1-\alpha})^{1-\rho}}{1-\rho} \quad (3)$$

where $0 \leq \alpha \leq 1$ is the weight on nondurable consumption in the composite good, and $\rho \geq 0$ is the inverse of the elasticity of intertemporal substitution for the composite good.

Budget Constraint. We write the standard intertemporal budget constraint for the household in terms of cash-on-hand. Households start any period t with a given amount of non-housing wealth, W_t , and receive an uncertain labor income, Y_t , that adds up to cash-on-hand, X_t . Given the amount of cash-on-hand, households decide how much to consume, C_t , to invest in housing, I_t at unit price Q_t , to repay on existing mortgage debt, ξ_t and how much new mortgage to take out, ϑ_t . Depending on the size of the housing investment, households also have to pay the corresponding adjustment cost, Θ_t , which we specify later.

$$X_{t+1} = R^X(X_t - C_t - Q_t I_t - \Theta_t - \xi_t + \vartheta_t) + Y_{t+1} \quad (4)$$

By deciding on the amount of consumption, housing investment, mortgage repayment and new mortgage take-out, households determine how much to save in a one-period bond, which yields a risk-free return, R^X . The non-housing wealth available in the next period is consequently the current period bond augmented by its rate of return, while the next period's cash-on-hand is the sum of the next period's non-housing wealth and next period's labor income.

Housing. In addition to the one-period bond, households have access to housing. Housing investment, I_t , adds to their existing housing, H_t . Hence the law of motion for the stock of housing is

$$H_{t+1} = H_t + I_t, \quad H_0 = H_{min} > 0 \quad (5)$$

The stock of housing is a continuum and bounded from below at $H_{min} > 0$. We do not model renting explicitly as our main focus is on the behavior of leveraged homeowners. Instead, we consider households with the minimum stock of housing to be renters. In the United Kingdom, the rental market is separate from the owner-occupied housing market: rental properties are, on average, significantly smaller than owned ones. For example, the 2012 English Housing Survey reports that 80% of properties smaller than 50 square meters are rental, whereas larger properties are predominantly owner-occupied.

Purchasing a home is usually associated with significant costs such as time spent looking for the preferred house or contractual costs such as legal fees. To capture these types of costs, we allow for adjustment costs $\Theta(I_t) \geq 0$, which we assume to depend on the value of housing investment

$$\Theta(I_t) = \delta Q_t I_t \quad (6)$$

Housing generates housing services, S_t , which yield instantaneous utility. There is no rental market in our model, so housing services can only be consumed by owning. We assume that there is a linear technology between housing and housing services.

$$S_t = bH_t \quad (7)$$

Housing can be used as collateral for mortgage loans. Households can get collateralized debt at a constant price of R^M up to a given fraction, $1 - \psi$, of the value of housing. So at the moment of the first mortgage take-out, the following inequality has to hold

$$M_t \leq (1 - \psi)Q_t H_t \quad (8)$$

where M_t is the mortgage and ψ can be interpreted as the mortgage down payment requirement.

Financial markets. We only allow households to have collateralized debt. There is no other type of debt available.¹⁴ Consequently households may have an incentive to take out a mortgage even if the mortgage rate is higher than the risk-free rate. Households with an existing mortgage, M_t can apply for a new mortgage, ϑ_t , but will have to keep repaying the existing mortgage, ξ_t . The law of motion for the mortgage stock is as

¹⁴Mortgages constitute the vast majority of loans in the household sector in the UK and the predominant investment market for British households is housing.

follows

$$M_{t+1} = R^M M_t + \vartheta_t - \xi_t \quad (9)$$

Next period's mortgage equals the existing mortgage with its interest, R^M , plus the new mortgage taken out minus the repayment on the existing mortgage¹⁵. We assume that repayment on the existing mortgage is bounded from below—households have to pay at least the interest on the value of the mortgage in each period. Also, there is a natural upper bound for repayment, which is paying back all the mortgage with all its interest.

$$r^M M_t \leq \xi_t \leq R^M M_t \quad (10)$$

As highlighted earlier, households are assumed to face a constraint on the level of mortgage they can take out. The maximum amount of mortgage they can have is a constant fraction of the value of their housing.

$$\vartheta_t \begin{cases} \leq (1 - \psi)Q_t H_t - R^M M_t & \text{if } (1 - \psi)Q_t H_t > R^M M_t \\ = 0 & \text{else} \end{cases} \quad (11)$$

It is important to see from equation (11), that the restriction on the mortgage relative to the value of housing is only enforced at the moment of taking out a new mortgage. As house prices fluctuate, the constraint can be violated for households with an existing mortgage as there is no mechanism through which households could insure themselves against this uncertainty. As a result, whenever the existing mortgage exceeds the maximum possible mortgage take-out, households cannot apply for a new mortgage. Whenever households borrow they are also subject to the terminal condition, $M_T = 0$, which prevents them from borrowing more than they can repay with certainty by the end of their life.

Uncertainty. In our baseline model households face uncertainty in two dimensions: idiosyncratic uncertainty over labor income and aggregate uncertainty over house prices. Following Zeldes (1989), labor income Y_t at any time before retirement is exogenously described by a combination of deterministic and random components

$$Y_t = Y_t^P Z_t \quad \log(Z_t) \sim N(-0.5\sigma_z^2, \sigma_z^2) \quad (12)$$

where Y_t^P is the permanent component and Z_t is the transitory component. Furthermore

¹⁵ R^M is the gross real mortgage rate, $R^M = 1 + r^M$

we assume that the permanent component can be described as

$$Y_t^P = G_t Y_{t-1}^P N_t \quad \log(N_t) \sim N(-0.5\sigma_n^2, \sigma_n^2) \quad (13)$$

with G_t being a deterministic function of age and N_t is the innovation. We also assume that the shocks (N_t and Z_t) are independent.

Labor income, Y_t , at any time after retirement is a constant fraction a of the last working year's permanent labor income. One can think of this as a pension that is wholly provided by the employer and/or the state.

$$Y_t = aY_W^P \quad (14)$$

The log of the house price is assumed to be determined by a random walk process with drift. In Section 4.3, we show that this assumption is consistent with the UK house price data.

$$\log Q_{t+1} = d_0 + \log Q_t + \log \varepsilon_t \quad \log(\varepsilon_t) \sim N(-0.5\sigma_\varepsilon^2, \sigma_\varepsilon^2) \quad (15)$$

The parametrization of uncertainty assumed in equations (12)-(13) and (15) applies in the baseline specification

Having all the details of the model specified, we can define the vector of state variables $\Omega_t = (X_t, H_t, M_t, Q_t, Y_t^P)$ and formulate the households's value function in period t in a recursive form as:

$$V_t(\Omega_t) = \max_{\{C_t, I_t, \xi_t, \vartheta_t\}} U(C_t, S_t) + \beta \mathbb{E}_t V_{t+1}(\Omega_{t+1}) \quad (16)$$

subject to the budget constraint, the income processes and the form of the utility function specified earlier.

4.2 Solution and Simulation

This life-cycle problem cannot be solved analytically, so we apply numerical techniques. Given the finite nature of the problem, a solution exists and can be obtained by approximating optimal policy functions by backward induction. We use the backward induction technique over the normalized value function of the households to obtain the optimal policy functions. Expectations in the model refer to uncertain incomes and house prices, while they are evaluated using the Gauss-Hermite approximation. Since shocks to incomes and prices, Z_{t+1} , N_{t+1} and ε_{t+1} are log-normally distributed random variables in each period, we are able to use a three-dimensional Gauss-Hermite quadrature to

approximate the expectations as follows

$$\begin{aligned}
\mathbb{E}_t V_{t+1}(\omega_{t+1}) &= \int V_{t+1}(\omega_{t+1}(Z, N, \varepsilon)) dF(Z)dF(N)dF(\varepsilon) \\
&= \int_{-\infty}^{\infty} \frac{1}{\pi} V_{t+1}(\omega_{t+1}(\sqrt{2}\sigma_Z Z, \sqrt{2}\sigma_N N, \sqrt{2}\sigma_\varepsilon \varepsilon)) e^{-(Z^2+N^2+\varepsilon^2)} \\
&\approx \sum_{i \otimes j \otimes k} \frac{1}{\pi} w_i^{GH} w_j^{GH} w_k^{GH} V_{t+1}(\omega_{t+1}(\sqrt{2}\sigma_Z Z_i^{GH}, \sqrt{2}\sigma_N N_j^{GH}, \sqrt{2}\sigma_\varepsilon \varepsilon_k^{GH}))
\end{aligned} \tag{17}$$

where ω_{t+1} is the vector of normalized state variables¹⁶; Z_i^{GH} , N_j^{GH} and ε_k^{GH} are the Gauss-Hermite nodes; and w_i^{GH} , w_j^{GH} and w_k^{GH} are the corresponding weights. Using backward induction, we get the optimal policy functions, which are shown in Figures B.2-B.4 in the Appendix.

Having calculated these policy functions, we simulate the behavior of households over their life-cycle. For each simulation we draw realizations for the two household-specific shocks, the permanent income shock and the transitory income shock; as well as realizations for the aggregate house price shocks which are identical for all the households. We disregard any kind of bequest motive: households start their life with zero wealth, and they receive only labor income. We define different groups of households, and simulate their behavior separately, taking into account that different cohorts experience aggregate shocks at different ages. Altogether we define 6 cohorts between ages 20 and 80 (10-year age interval each) and simulate 1000 households in each cohort.

Consider next how we amend the baseline model to take account of the four channels whereby the 2009 recession may have affected household behavior, as discussed in the introduction.

Permanent Income. A recession may occur when all households face a negative income shock at the same time. Indeed, data from the ONS show that incomes for all UK households fell by 5-10% in the crisis, relative to trend growth.¹⁷ We model this scenario by assuming that the expected value of the permanent income shock, N_t , decreases from 1 to 0.9 for one period in 2009. Hence, the level of the income shock becomes

$$N_{2009} = 0.9N_{2008}$$

for each household and the rest of the shocks are drawn from the same lognormal dis-

¹⁶Following Carroll (1992), variables are normalised by permanent income for ease of computation.

¹⁷We calculate this using total UK compensation of employees from the Quarterly National Accounts. By 2011, incomes for all households fell by over 10% from their pre-crisis peak.

tribution we have introduced in Section 4.1,

$$\log(N_t) \sim N(-0.5\sigma_n^2, \sigma_n^2).$$

Note that this aggregate, negative permanent income shock reduces income not only in 2009, but after the crisis as well.

Variance of Permanent Income. The other alternative cause of a recession is that from 2009 onward, households experience a higher level of uncertainty around their future labor income. We simulate this situation by considering an increase in the variance of the permanent shocks to income. Following the calculations of Blundell, Low, and Preston (2013) based on British data¹⁸, we assume that the variance of the permanent income shock after 2008 becomes three times higher than before and it permanently stays there.

$$\begin{aligned} \log(N_t) &\sim N(-0.5\sigma_n^2, \sigma_n^2) && \text{if } t < 2009 \\ \log(N_t) &\sim N(-1.5\sigma_n^2, 3\sigma_n^2) && \text{if } t \geq 2009 \end{aligned}$$

Note that in order to keep the expected value of the level of the shock fixed, $\mathbb{E}(N_t) = 1$, whenever we change the variance of the lognormal distribution we change the mean as well.

Down Payment Requirement. We can think of the recession as a period with a tightening credit market. In fact, banks changed their credit conditions on new borrowing dramatically after 2008. Instead of an average 10% pre-crisis down payment requirement for example, they demanded a 15 – 20% down payment for new mortgages after the crisis. In our model, down payment requirement is led by a single parameter, ψ , as seen in equation (8). Hence we set this parameter accordingly:

$$\begin{aligned} \psi &= 0.1 && \text{if } t < 2009 \\ \psi &= 0.15 && \text{if } t \geq 2009 \end{aligned}$$

House Prices. The last alternative we consider is the collapse of house prices after 2008. Instead of our estimated random walk process for the house prices reported in

¹⁸In the empirical part of their paper, the authors use data from the Family Expenditure Survey (FES).

Table 7, we feed the observed house price data, shown in Figure B.1 in the Appendix, into the model.

4.3 Model Calibration

Most of the parameter values for the life-cycle model presented above are adopted from the existing literature. The benchmark model parameters are collected in Table B.1 in the Appendix. Here, we discuss only the values of parameters which we estimate directly.

Income process. To obtain the age-specific component of the life-cycle income profiles (G), we fit a second-order age polynomial to the logarithm of cohort income data gathered from the Family Expenditure Survey between 1992 and 2012.

$$\ln y_t^c = \gamma_0 + \gamma_1 age_{c,t} + \gamma_2 \frac{age_{c,t}^2}{10} \quad (18)$$

where c refers to cohort averages. The regression results are presented in Table 6.

| | $\log y^c$ |
|-------------------------|----------------------|
| <i>Age</i> | 0.110*** (0.003) |
| <i>Age</i> ² | -0.012*** (0.000) |
| Constant | 7.943*** (0.058) |
| Observations | 945 |
| R-squared | 0.614 |

Standard errors are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 6: Income process

House price process. We estimate the house price process based on Nationwide's House Price Index adjusted for retail prices (using the Office for National Statistics Retail Price Index) for quarters 1992q1-2012q2. We estimate an $AR(1)$ process with

linear trend for the logarithm of the real house price index.

$$\log Q_{t+1} = q_0 + q_1 t + \rho_h \log Q_t + \log \varepsilon_t \quad (19)$$

The result of the estimation is in Table 7. The persistence parameter ρ_h of the log real house price is estimated to be very close to one and the unit root tests do not reject the null hypothesis that this parameter is actually 1. It implies that the log real house price process can be approximated by a random walk. The estimated quarterly variance of the house price shock, σ_ε^2 , equals 0.0006, which corresponds to 0.0025 at an annual frequency.

| | log Q | |
|------------------------|---------------------|--------------------|
| t | -0.0003 (0.0002) | |
| log Q (-1) | 1.010*** (0.017) | 1 constrained |
| Constant | -0.088 (0.179) | 0.006** (0.002) |
| σ_ε^2 | 0.0007 | 0.0007 |
| Observations | 84 | 84 |
| R-squared | 0.995 | |

Standard errors are in parentheses. *** $p < 0.01$,
** $p < 0.05$, * $p < 0.1$

Table 7: House price process

5 Simulation Results

In this section, we present results from our simulation exercise for households aged between 20 and 60. In our baseline simulation, we assume no recession in 2009 and a random walk process for real house prices as estimated in Table 7. Then, following Alan, Crossley, and Low (2015), we model the 2009 recession in four different ways: as a negative shock to income, an increased uncertainty around future income, a reduction in the supply of credit, and a fall in house prices (i.e., a reduction in housing wealth). Our aim is to show whether different types of recessions can generate the three empirical facts about consumption behavior we demonstrated in Sections 2 and 3: that both low and high-leveraged households cut their consumption significantly in response to the crisis; that high-leveraged households cut their consumption by much more than low-leveraged

households; and that this difference is mostly driven by young households.

We do not aim for a quantitative match of the empirical and simulated moments, but, as we demonstrated in Section 4.3, we calibrate the model to be as realistic as possible. In Table 8, we report some basic average statistics of households' pre-crisis (2006/2007) leverage and wealth status both from the data and the simulated model. Without targeting the empirical moments, the model gets close to what we observe in the data. Mortgages in the model, though, play a more important role than in reality (see the LTI and LTV ratios), which is the result of the simplified mortgage market in our model that allows households to take-out additional mortgages at no cost.

| | Data | Model |
|--------------------------|-------------|--------------|
| Loan-to-income (LTI) | 2.4 | 2.8 |
| Loan-to-Value (LTV, %) | 39 | 59 |
| Housing Wealth to Income | 4.4 | 4.8 |
| Total Wealth to Income | 4.8 | 5.2 |

Table 8: Comparison of means in 2006/2007

Similarly to the empirical part before, our next step is to define subgroups of households along two dimensions: leverage status and age. We then report our simulation results for these subgroups. High-leveraged households are defined as those with an above-mean pre-crisis LTI¹⁹, whereas low-leveraged households are those with a below-mean LTI. In our model, as shown in Table 8, mean LTI is 2.8. Young households are defined as those between 20 and 40 years of age, while old households are the ones between 40 and 60 years of age.

5.1 Results by Leverage

Figure 5 summarizes the findings from our simulations that test the four different channels we described earlier: the permanent income shock; the uncertainty shock; the credit supply shock; and the house price shock. Each graph shows the difference between consumption growth simulated under one particular shock and the consumption growth simulated under the baseline model, i.e. no shock in 2009. The solid line in each graph

¹⁹Note that we could have chosen the median LTI, but it is very close to the mean value. In 2006/2007 the median was 2.84.

corresponds to high leveraged households, whilst the dashed line corresponds to low leveraged ones.

There are several things to notice here. First, all of these shocks lead to a substantial fall in consumption growth relative to the **baseline**, ‘no recession’ scenario, for both types of households. There is only one exception: low-leveraged households who face an increase in income uncertainty barely decrease their consumption, and they only do so in 2010. The biggest relative fall in consumption growth from 2008 to 2009 is associated with high-leveraged households after an increase in the down payment requirement (one can see about a 15-percentage-point gap in consumption growth relative to the baseline). The drop in permanent income implies a fall in consumption growth for high-leveraged households of similar size, while the other two channels lead to much smaller consumption responses.




Figure 5: Deviation of consumption growth from no recession

Second, consumption growth of high- and low-leveraged households (relative to the baseline scenario) under each of the four channels are very different. In particular, the permanent income shock implies the highest, about a 7 percentage point difference in the growth rate of consumption of high- and low-leveraged households. The shock to down payment results in a difference of 3.5 percentage points, while increased uncertainty re-

sults in a difference of 2.5 percentage points. The house price shock, however, implies no difference between consumption growth of high- and low-leveraged households.

If we compare these results to our empirical observations in Section 3, we conclude that the changes in both income variance and house prices would themselves not be enough to explain household's consumption behavior.

More specifically, an increase in the variance of permanent income fails to explain our first empirical observation, i.e. that low-leveraged households also decrease their consumption growth as seen in Figure 2. In a textbook version of the life-cycle model, an increase in the uncertainty around income would affect savings positively (as discussed in Carroll (1992) for example) leading to a decrease in consumption (or to a negative income growth). Our model is different because households also have access to the credit market through housing. Low-leveraged households, for example, can top-up their mortgages easily if necessary, which means that they can secure a smooth consumption profile in the presence of higher income uncertainty without saving more. For this reason, low-leveraged households in our model only change their consumption behavior marginally after an increase in income uncertainty.

Changes in house prices also fail to explain our second empirical observation, i.e. that high-leveraged households decreased their consumption growth by more than low-leveraged households. In particular, high- and low-leveraged households' consumption growth is almost identical in 2009. Independent of mortgage status, a house price drop generates a permanent reduction in a household's wealth. As a result, both types of households feel less wealthy and they re-optimize their consumption plans, which includes an immediate, similar-sized drop in their consumption.

On the contrary, and as seen in Figure 5, a permanent income shock or a down payment shock is able to explain our first two empirical observations. Each of these shocks generate sharp drops in consumption growth both for high- and low-leveraged households, while leading high-leveraged households to decrease their consumption growth more than low-leveraged ones. For this reason, we solely concentrate on these two channels in the remainder of the paper.

5.2 Results by Age

In this section, we extend our analysis of the simulation results in the same direction as we did in our empirical analysis by examining consumption behavior by age group. In doing so, our aim is to explore whether a permanent income or a down payment

shock can also explain our third empirical fact: that the observed consumption growth difference of high- and low-leveraged households is mainly driven by young households.

In Figure 6 below, we present a more disaggregated version of Figure 5 for permanent income and down payment shocks. Here, high- and low-leveraged households are divided into two additional subgroups based on their age. This figure can be interpreted the same way as in Figure 5 above. Both graphs plot the difference between the consumption growth simulated under one particular shock and the consumption growth simulated in the baseline setting of no shock in 2009. The solid lines in each graph correspond to high leveraged households, whilst the dashed lines correspond to low-leveraged ones.

Permanent Income. As Figure 6 shows, a negative shock to households' permanent income leads to a larger drop in consumption growth for the young compared to the old. In 2009, the average consumption growth of high-leveraged young households is about -12%, while it is about -4% for the low-leveraged young ones. At the same time, the average consumption growth of high-leveraged old households is about -10%, while that of the low-leveraged old is about -2%.

The reason why young households react to a permanent income drop more than old ones can be explained by comparing their lifetime wealth. When a negative permanent income shock hits, it affects the level of household income permanently. Consequently, the younger the household is at the time of the income shock, the greater the negative wealth effect that is induced by the permanently lower levels of income. Therefore, a same-sized unexpected income shock makes young households adjust their consumption by more than old ones.

Figure 6 also reveals a sizeable difference between the consumption growth of high- and low-leveraged young households. Highly leveraged young households decrease their consumption growth by around 8 percentage points more than the low-leveraged young. Moreover, the consumption growth difference between high- and low-leveraged old households is almost identical to that of young households - which is again about 8 percentage points.

As a result, considering the recession only as a drop in households' permanent income cannot explain our third empirical fact: that the observed consumption growth difference of high- and low-leveraged households is driven by young households.

Down payment. As Figure 6 shows, a negative shock to credit supply leads to a larger drop in consumption growth of old households relative to young ones. In 2009, the average consumption growth of high-leveraged old households is about -16%, while that of the low-leveraged old is roughly -14%. At the same time, the average consumption

growth of the high-leveraged young households is about -14%, while that of the low-leveraged old is about -9.5%.

The reason why old households are more responsive to an increase in down payment requirements than young households can be explained by their need to top up mortgages. Our policy functions in Figure B.4 show that old households are more likely, on average, to take out an additional mortgage than the young. One way to think about it in our model framework is to consider home downsizing: in order to access home equity, old households might find it less costly to top-up their mortgages rather than to downsize, which comes with sizeable housing transaction costs. Topping up eventually has the same effect as downsizing, which decreases the value of households' home equity. But in contrast to downsizing, topping up has the benefit of keeping the housing service flow high.²⁰ Old households' higher demand for topping up their mortgages induced by their need for downsizing leads them to react more to changes in credit supply conditions.

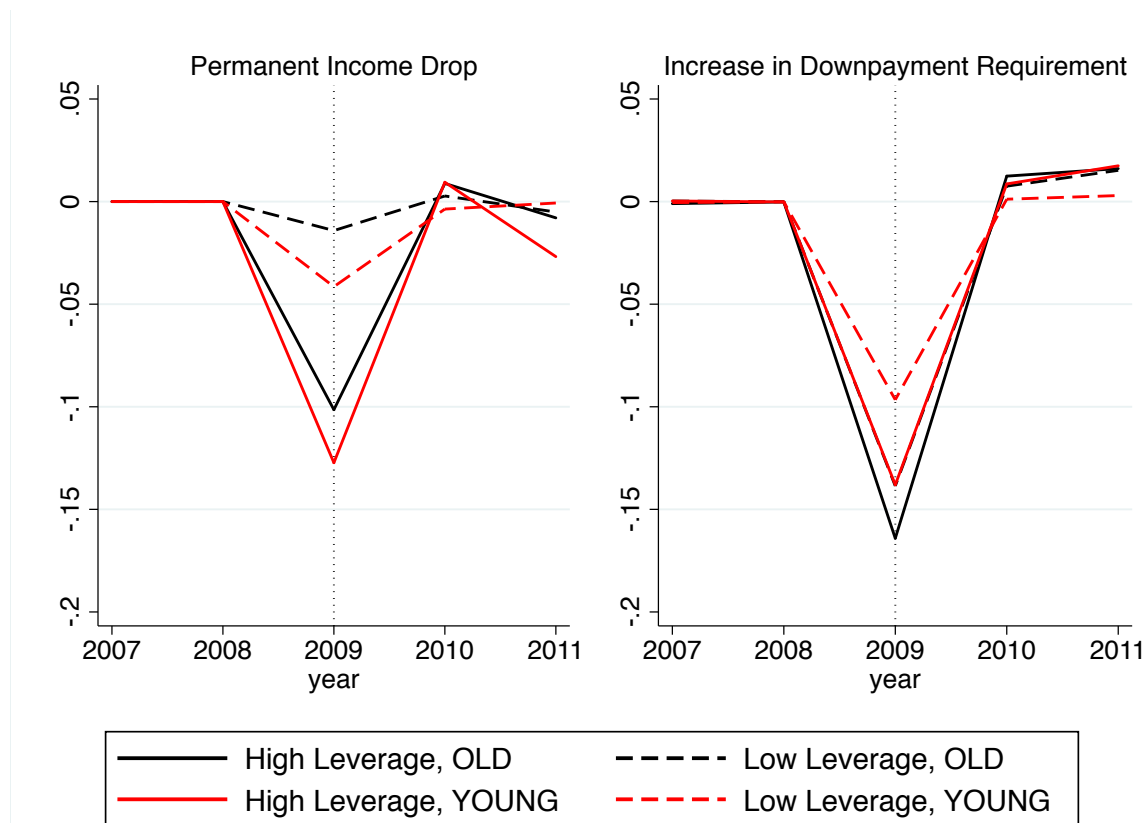


Figure 6: Deviation in consumption growth from no recession

Figure 6 also reveals a sizeable difference between the consumption growth of high- and low-leveraged young households. High-leveraged young households decrease their consumption growth by around 4 percentage points more than low-leveraged young households.

²⁰As households don't change the actual size of their house, but their home equity only.

holds. After an unexpected increase in the down payment requirement, high-leveraged young households, who are likely to be liquidity constrained at the same time, may find it difficult to withdraw additional home equity to finance consumption. This could force them to decrease their consumption by more than low-leveraged young households, who are able to top up their mortgages even under the stricter credit conditions.

The gap between the consumption growth of high- and low-leveraged old households is almost half the size of the gap in consumption growth of high- and low-leveraged young households.²¹ High-leveraged old households decrease their consumption growth by only around 2 percentage points more than low-leveraged old households. Our policy functions in Figure B.4 show that new mortgage take-outs of old households vary a lot by their leverage status. High-leveraged old households, on average, typically take out less additional mortgage than low-leveraged ones. Consequently, an increase in the down payment requirement could affect high- and low-leveraged old households similarly by pushing them towards their credit limits. High-leveraged households get to their credit limits because of their initially high mortgage levels, whilst low-leveraged ones reach their credit limits by taking-out higher levels of additional mortgages. As a result, both groups may decrease their consumption growth significantly after the tightening of down payment requirements.

As such, considering the recession only as a drop in credit supply can explain our third empirical fact. Indeed, the difference in consumption growth between high- and low-leveraged young households is 4 percentage points, double that of old households. Our model with a credit supply shock implies that the observed consumption growth difference of high- and low-leveraged households is driven by young households.

5.3 Evaluation of the Results

In Section 3, we showed three important empirical findings. First, both high- and low-leveraged households cut their consumption in response to the crisis. Second, high-leveraged households cut their consumption by more than low-leveraged ones. Third, the sizeable differences in the consumption drop between high- and low-leveraged households were driven by the young. After demonstrating these empirical facts, in Sections 4 and 5 we modelled the recession in four different ways and presented the corresponding results of our simulated economy, which helped us understand the most likely causes behind the observed consumption differences.

²¹Note that the consumption growth line of low-leveraged old households, the dashed black line, coincides with the line of high-leveraged young households, the solid red line.

We showed that both a drop in the permanent income or an increase in the down payment requirement could potentially explain the first two empirical facts. However, only an increase in the down payment requirement is able to match the third empirical fact: the increased down payment requirement makes high-leveraged young households more constrained than low-leveraged young ones, which leads them to decrease their consumption significantly more than low-leveraged households. The difference between high- and low-leveraged consumption growth for young households was as high as around 4 percentage points. This difference between high- and low-leveraged old households was less pronounced, at about 2 percentage points.

The empirical findings together with our simulation results makes us conclude that the most important driver of the empirical co-movement between household leverage and consumption growth after 2008 is the reduction in credit supply. The main mechanism is that after an unexpected increase in the down payment requirement, high-leveraged young households, who are likely to be liquidity constrained at the same time, find it difficult to withdraw additional equity from their housing. This leads them to decrease their consumption by more than low-leveraged young households, who are able to top up their mortgages even under the stricter credit conditions.

It is important to note that the basis of our analysis was a one-by-one examination of individual shocks. Clearly, one might want to consider the combination of these individual shocks in order to evaluate the quantitative effects of the financial crisis. Our goal however was different in this paper. We wanted to determine the main driver behind the substantial consumption growth difference between high- and low-leveraged households in 2009, as seen in Figure 2.

6 Conclusion

In this paper we investigate households' consumption behavior around the 2008 financial crisis in the UK. We make two important contributions to the existing literature.

We first document the existence of a relationship between the financial crisis and differences in households' consumption responses based on leverage status and age. High-leveraged households made much larger cuts to their consumption after 2008 compared to low-leveraged ones. On average, a one unit increase in leverage, measured by the loan-to-income ratio, is correlated with a 3-percentage point drop in consumption growth. On the other hand, the correlation is twice as strong for young households than for old. A one unit increase in leverage led to a 4-percentage point drop in consumption growth for

young households, but only to a 2-percentage point drop in consumption growth for old households. This suggests that the sizeable drop of the aggregate consumption growth, seen in Figure 2, was likely driven by young households.

After our detailed empirical analysis, we then simulate households' consumption behavior in a structural model. We show that the observed differences in consumption growth based on leverage status and age after 2008 are best explained by an exogenous reduction in credit supply. Besides the negative credit market shock, we consider three counterfactual channels by which a recession can affect consumption behavior: a fall in the level of permanent income, rising uncertainty around future income and a decline in house prices. None of these other channels can match the observed consumption patterns.

The results from this paper are very suggestive, but more work is required for a better quantitative match of our empirical results. The most important extension of this work is to introduce discrete housing choice to make it possible to explain lifetime borrowing profiles better. In future work, we plan to explore this.

A Appendix

A.1 Summary Statistics

Table A.1 reports summary statistics from the FES and WAS based on cross-sectional data.

A.2 Robustness checks for Data Exercise

A.3 Loan-to-value

As explained above, our preferred synthetic panel definition is one split by single date of birth years and mortgagor status. Below we show that the coefficients on LTV's are also not affected by choice of cohort definition chosen, similar to what we obtain for LTI's in Table 4.

A.4 Selection

Our preferred empirical equation specification uses mortgage status to define cohorts. As buying a house and taking out a mortgage are choice variables, there may be selection using mortgage status in cohort definition. Moffitt (1993) suggests that one can use deterministic variables to predict selection into the choice variable. We can then use predicted status, rather than actual status, to group households into mortgagors and non-mortgagors. An application of this approach is found in Attanasio, Banks, and Tanner (2002), who use exogenous or time-invariant observables to predict shareholder status. We follow a similar methodology to predict mortgagor status. But an additional complication arising from our analysis is that we also need to impute mortgage levels too. We estimate a tobit model to simultaneously predict mortgagor status and mortgage debt levels at the household level, in the cross-sectional data. Following that, we group cohorts into date of birth and predicted mortgagor status as described using our preferred cohort definition.

$$debt_h^* = \kappa_0 + \kappa_1 age_h + \kappa_2 educ_h + \kappa_3 age * educ_h + e_h \quad (A.1)$$

Where $debt_h^*$ is the latent variable of outstanding level of mortgage debt for household

| Cohort definition | Year of birth | Mortgagor / Non-Mortgagor |
|------------------------------|---------------------------|------------------------------|
| | [1] | [2] |
| Dependent var: ΔC_t | | |
| LTV_{t-1} | -0.128* (0.064) | -0.142*** (0.039) |
| ΔY_t | 0.743*** (0.124) | 0.551*** (0.117) |
| $\Delta Housing\ wealth_t$ | 0.123 (0.096) | 0.074** (0.035) |
| $\Delta Financial\ wealth_t$ | 0.004 (0.020) | 0.005 (0.023) |
| Constant | -0.011 (0.029) | -0.028** (0.012) |
| Controls | Yes | Yes |
| Observations | 45 | 77 |

Robust standard errors in parentheses, *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. ΔC_t is the change in log non-housing consumption in 2009/2010, denoted by t . ΔY_t is the change in after-tax income, net of mortgage interest payments at t . LTV_{t-1} is the ratio of mortgage debt to housing wealth in 2006/2007, denoted by $t - 1$. $\Delta Housing\ wealth_t$ and $\Delta Financial\ wealth_t$ are changes in housing wealth and total financial wealth respectively. Controls are the change in the number of adults and children. Each column represents a synthetic panel regression from 2 pooled years, constructed using either [1] year of birth, or [2] year of birth, split by mortgage status.

Table A.2: Baseline results (Loan-to-Value)

h in 2006/2007; age is a vector of age polynomials for the head of households²²; $educ$ is a vector of indicator variables for whether the head of the household has a high-school degree, college, or a university degree

A further complication that arises from this method is we are unable to predict wealth for the same set of individuals, given that this variable is spliced using a secondary data source (WAS). This is of particular importance because we would also need to predict housing wealth for those who are predicted to have mortgage debt. We therefore opt to exclude wealth from the re-estimation of Equation 1, as below:

$$\Delta c_{it} = \phi_0 + \phi_1 L\hat{T}I_{it-1} + \phi_2 y_{it} + \phi_3 x_{it} + e_{it} \quad (\text{A.2})$$

Where \hat{i} is the cohort predicted to have mortgage debt and $L\hat{T}I$ is the predicted debt to income ratio.

Table A.3 compares results of actual to predicted mortgage status. As mentioned above, the WAS data show that 90% of households who had a mortgage in 2006/2007 also had one in 2009/2010. Indeed, the tobit equation correctly classifies 70% of households who have a mortgage to also have a predicted mortgage in 2006/2007. It is therefore unsurprising that the LTI coefficients from both regressions in Table A.3 are very similar.

There are nevertheless two drawbacks from using predicted mortgage status as our baseline estimation. First, we are unable to control for housing wealth, which makes the LTI coefficient endogenous. Second, there are some age groups where the model predicts households to be mostly mortgagors and others where they are mostly non-mortgagors; and as we retain the rule that there must be a sufficient number of household-level data points in each panel observation, we end up with fewer observations when using predicted status. This means imprecise estimates, leading to inefficiency.

²² age, age^2, \dots, age^5

| | Actual | Predicted |
|------------------------------|-----------------------------|---------------------------|
| | [1] | [2] |
| Dependent var: ΔC_t | | |
| LTI_{t-1} | -0.026*** (0.008) | |
| $\hat{LTI}_{(t-1)}$ | | -0.030* (0.017) |
| ΔY_t | 0.546*** (0.117) | 0.755*** (0.100) |
| $\Delta Housing\ wealth_t$ | 0.073** (0.035) | - |
| $\Delta Financial\ wealth_t$ | 0.004 (0.023) | - |
| Constant | -0.028** (0.012) | -0.045*** (0.015) |
| Controls | Yes | Yes |
| Observations | 77 | 48 |

Robust standard errors in parentheses, *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. ΔC_t is the change in log non-housing consumption in 2009/2010, denoted by t . ΔY_t is the change in after-tax income, net of mortgage interest payments at t . LTI_{t-1} is the ratio of mortgage debt to annual income in 2006/2007, denoted by $t - 1$. $\Delta Housing\ wealth_t$ and $\Delta Financial\ wealth_t$ are changes in housing wealth and total financial wealth respectively. Controls are the change in the number of adults and children. Regressions are from a synthetic panel using 2 pooled years, where the panel is constructed using year of birth and mortgage status [1], or predicted mortgagor status [2]. Predicted debt and mortgagor status are calculated using regressions equations from A.1 and A.2.

Table A.3: Using predicted mortgagor and debt

A.5 Extending the time-period

While the main focus of the paper is on the period immediately after the financial crisis, we check for persistence by extending the estimation period to 2011/2012. In other words, we extend our estimation period to cover the change in consumption from 2006/2007 out to 2011/2012, rather than 2009/2010. The coefficients on LTI and LTV continue to remain statistically significant but also become more negative (Table A.4). This suggests that the larger cuts in spending by indebted households had not dissipated, at least up to 2012.

A.6 LTI bands

Chart 2 above shows that highly leveraged households made a bigger adjustment to consumption than low-leveraged households. To test whether this effect is larger as leverage increases, we create a flag for cohorts who had a loan-to-income greater than k in 2006/2007, where k is 2, 2.5 or 3.

$$\Delta c_{it} = \gamma_0 + \gamma_1(LTI > k)_{it-1} + \gamma_2\Delta y_{it} + \gamma_3\Delta w_{it} + \gamma_4x_{it} + e_{it} \quad (\text{A.3})$$

Table A.5 reports these results. The Table shows that for higher the loan-to-income threshold, the coefficients become larger in magnitude. Moreover, the effects appear non-linear— the coefficient of -0.049 more than doubles when moving from an $LTI > 2$ to an $LTI > 3$. These numbers are consistent with previous analysis showing a correlation between households with high pre-crisis loan-to-income ratios and a subsequent fall in consumption growth after the crisis— and that the effect may be non-linear.

| Time period, t : | 2009/2010 | 2009/2010 | 2011/2012 | 2011/2012 |
|----------------------------------|-----------------------------|-----------------------------|-----------------------------|-----------------------------|
| | [1] | [2] | [3] | [4] |
| Dependent var: ΔC_t | | | | |
| LTI_t | -0.026*** (0.008) | | -0.030*** (0.007) | |
| LTV_t | | -0.142*** (0.039) | | -0.166*** (0.037) |
| ΔY_t | 0.546*** (0.117) | 0.551*** (0.117) | 0.572*** (0.101) | 0.569*** (0.102) |
| $\Delta Housing\ wealth_{(t)}$ | 0.073** (0.035) | 0.074** (0.035) | 0.015 (0.032) | 0.016 (0.031) |
| $\Delta Financial\ wealth_{(t)}$ | 0.004 (0.023) | 0.005 (0.023) | -0.015 (0.022) | -0.015 (0.021) |
| Constant | -0.028** (0.012) | -0.028** (0.012) | -0.032** (0.013) | -0.031** (0.013) |
| Controls | Yes | Yes | Yes | Yes |
| Observations | 77 | 77 | 74 | 74 |

Robust standard errors in parentheses, *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Each column represents a regression for a different time period. ΔC_t is the change in log non-housing consumption in t ; where t is 2009/2010 [column 1 & 3], or 2011/12 [column 2 & 4] respectively. ΔY_t is the change in after-tax income, net of mortgage interest payments at t . LTI_{t-1} is the ratio of mortgage debt to annual income in 2006/2007, denoted by $t - 1$. $\Delta Housing\ wealth_t$ and $\Delta Financial\ wealth_t$ are changes in housing wealth and total financial wealth respectively. Controls are the change in the number of adults and children. Regressions are from a synthetic panel using 2 pooled years, where the panel is constructed using year of birth and mortgage status.

Table A.4: Longer time-period

| | LTI>2 | LTI>2.5 | LTI>3 |
|------------------------------|----------------------------|----------------------------|-----------------------------|
| | [1] | [2] | [3] |
| Dependent var: ΔC_t | | | |
| $LTI_t > 2$ | -0.049** (0.020) | | |
| $LTI_t > 2.5$ | | -0.065** (0.027) | |
| $LTI_t > 3$ | | | -0.132*** (0.033) |
| ΔY_t | 0.529*** (0.118) | 0.546*** (0.118) | 0.566*** (0.116) |
| $\Delta Housing\ wealth_t$ | 0.086** (0.035) | 0.089** (0.034) | 0.091*** (0.034) |
| $\Delta Financial\ wealth_t$ | 0.008 (0.025) | 0.006 (0.024) | -0.001 (0.019) |
| Constant | -0.038*** (0.012) | -0.039*** (0.011) | -0.042*** (0.010) |
| Observations | 77 | 77 | 77 |

Table A.5: LTI Bands

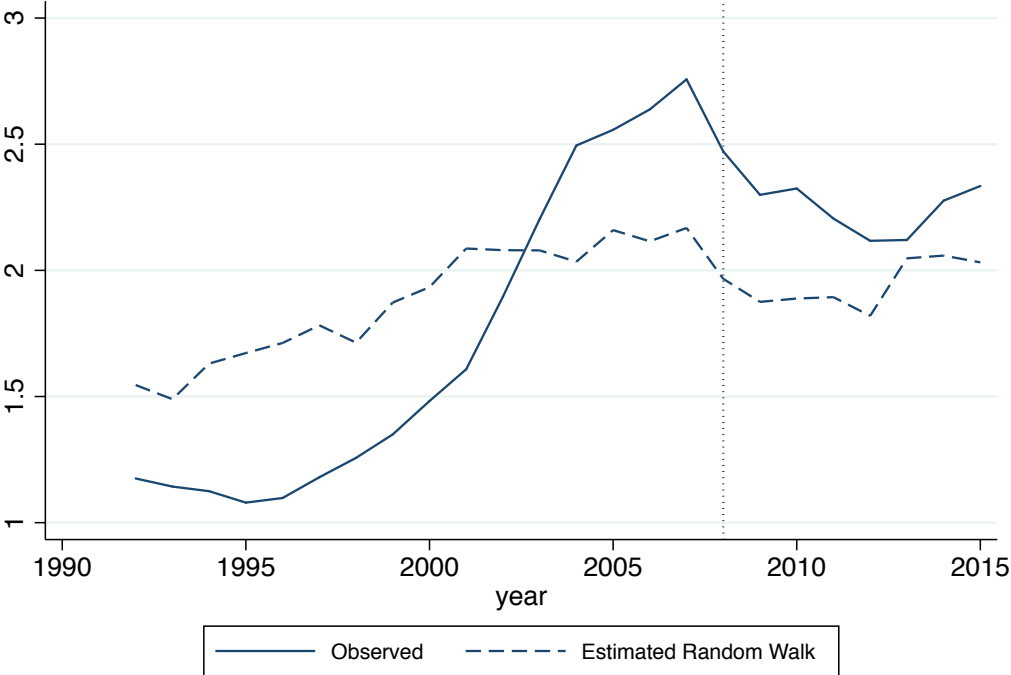
Robust standard errors in parentheses, *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. ΔC_t is the change in log non-housing consumption in 2009/2010, denoted by t . ΔY_t is the change in after-tax income, net of mortgage interest payments at t . $LTI_{t-1} > 2, 2.5$, and 3 are variables for the proportion of households in each cohort that have an outstanding mortgage to income ratio of more than 2, 2.5 or 3 respectively in 2006/2007, denoted by $t - 1$. $\Delta Housing\ wealth_t$ and $\Delta Financial\ wealth_t$ are changes in housing wealth and total financial wealth respectively. Controls are the change in the number of adults and children. Regressions are from a synthetic panel using 2 pooled years, where the panel is constructed using year of birth and mortgage status.

B Simulation

B.1 Model Parameters

The model parameters are presented in Table B.1.

B.2 Real House Prices



Nationwide's Calculations using Retail Prices to adjust House Prices

Figure B.1: Real house price evolution

| Parameter | Value | Source |
|----------------------|--------|---------------------------------------|
| T | 60 | |
| W | 45 | |
| β | 0.95 | |
| ρ | 1.5 | Blundell, Browning, and Meghir (1994) |
| Constant | 7.943 | Own calculations, FES |
| Age | 0.110 | Own calculations, FES |
| Age ² /10 | -0.012 | Own calculations, FES |
| a | 0.72 | Own calculations, FES |
| σ_n | 0.147 | Carroll and Samwick (1997) |
| σ_z | 0.21 | Carroll and Samwick (1997) |
| σ_ε | 0.025 | Own calculation, Nationwide |
| ρ_h | 1 | Own calculation, Nationwide |
| R^X | 1.02 | Gourinchas and Parker (2002) |
| R^M | 1.03 | |
| α | 0.83 | Kovacs (2016) |
| b | 0.06 | Kovacs (2016) |
| δ | 0.05 | |
| ψ | 0.2 | |

Table B.1: Parameters for the benchmark model

B.3 Policy Functions

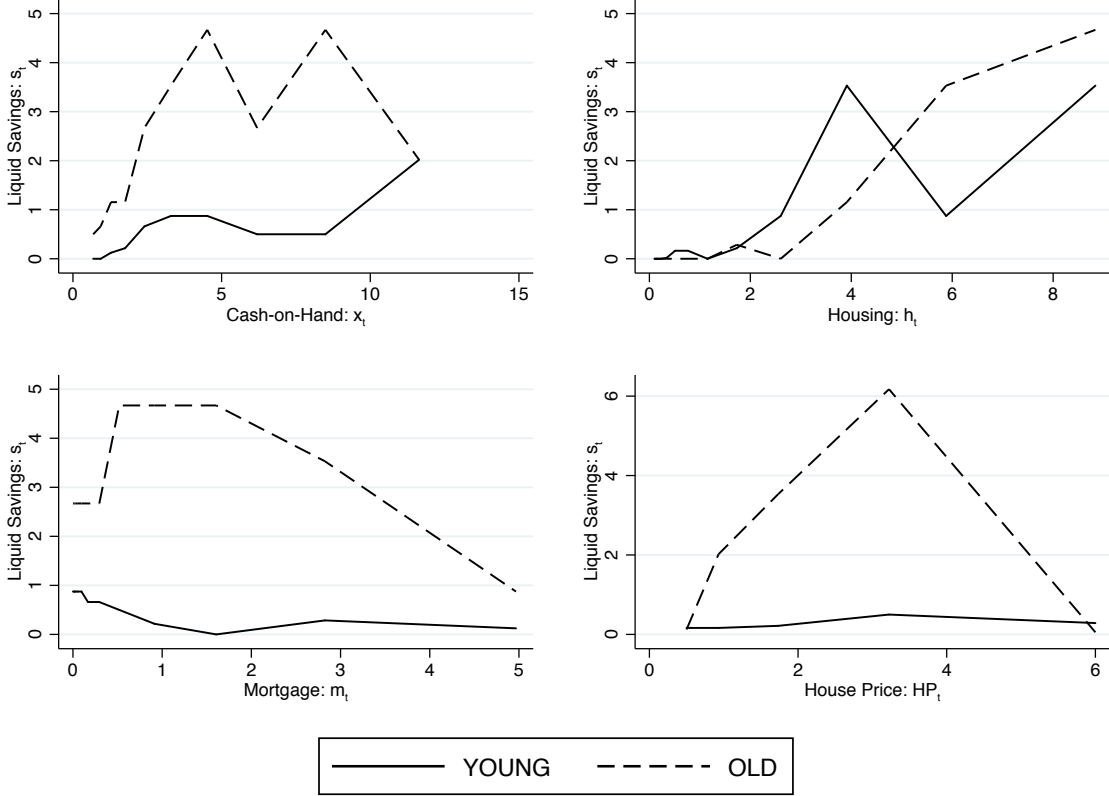


Figure B.2: Policy Functions for Liquid Savings




Figure B.3: Policy Functions for Housing Investment




Figure B.4: Policy Functions for New Mortgage

B.4 Baseline Simulation Results

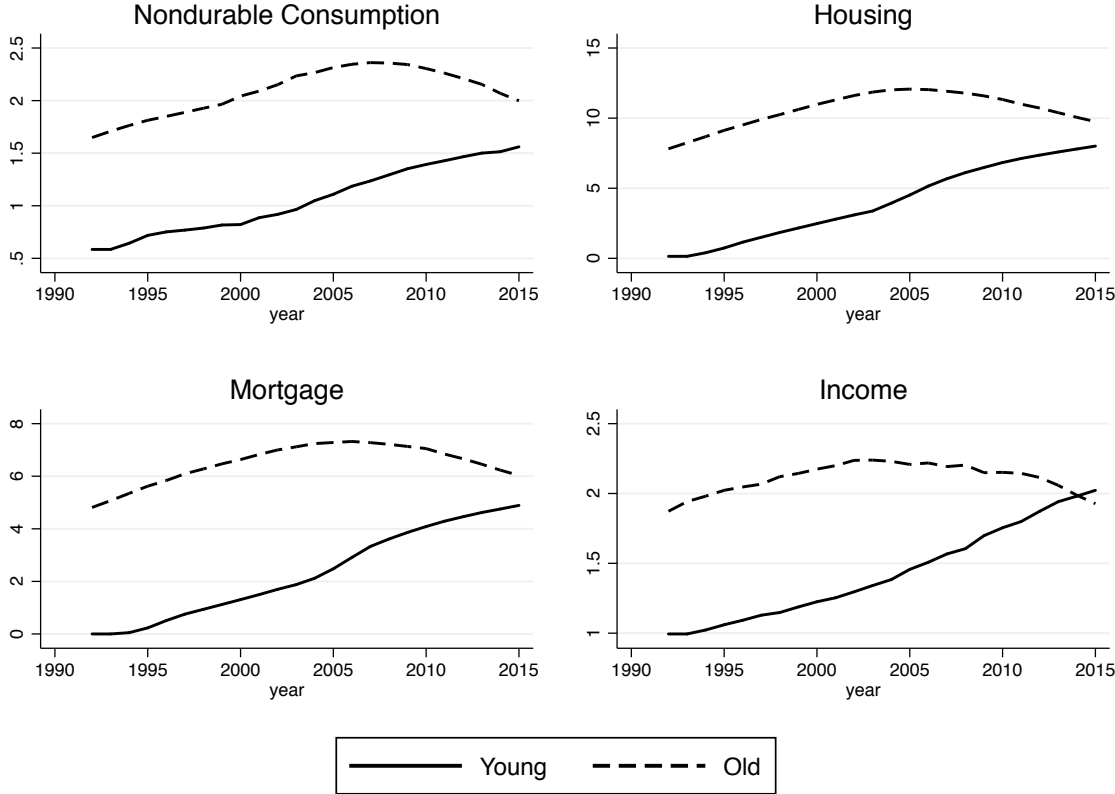


Figure B.5: Baseline Results

References

- Alan, Sule, Thomas Crossley, and Hamish Low. 2015. "Saving on a rainy day, borrowing for a rainy day." *Institute for Fiscal Studies Working Paper* .
- Andersen, Asger Lau, Charlotte Duus, and Thais Lærkholm Jensen. 2014. "Household debt and consumption during the financial crisis: evidence from Danish micro data." Tech. rep., Danmarks Nationalbank Working Papers.
- Attanasio, Orazio P., James Banks, Costas Meghir, and Guglielmo Weber. 1999. "Humps and Bumps in Life-Time Consumption." *Journal of Business and Economics Statistics* 17(1):22–35.
- Attanasio, Orazio P, James Banks, and Sarah Tanner. 2002. "Asset Holding and Consumption Volatility." *Journal of Political Economy* 110 (4):771–792.
- Attanasio, Orazio P, Laura Blow, Robert Hamilton, and Andrew Leicester. 2009. "Booms and busts: Consumption, house prices and expectations." *Economica* 76 (301):20–50.
- Attanasio, Orazio P., Renata Bottazzi, Hamish Low, Lars Neisham, and Matt Wakefield. 2012. "Modeling the Demand for Housing over the Lifecycle." *Review of Economic Dynamics* 15 (1):10–53.
- Attanasio, Orazio P. and Martin Browning. 1995. "Consumption over the Life Cycle and over the Business Cycle." *American Economic Review* 85(5):1118–1137.
- Attanasio, Orazio P., Andrew Leicester, and Matthew Wakefield. 2011. "Do House Prices Drive Consumption Growth? The Coincident Cycles of House Prices and Consumption in the UK." *Journal of European Economic Association* 9 (3):399–435.
- Attanasio, Orazio P. and Susann Rohwedder. 2003. "Pension Wealth and Household Saving: Evidence from Pension Reforms in the United Kingdom." *American Economic Review* 93(5):1499–1521.
- Attanasio, Orazio P. and Guglielmo Weber. 1993. "Consumption Growth, the Interest Rate and Aggregation." *Review of Economic Studies* 60(3):631–649.
- Attanasio, Orazio P and Guglielmo Weber. 1994. "The UK consumption boom of the late 1980s: aggregate implications of microeconomic evidence." *Economic Journal* :1269–1302.

- Attanasio, Orazio P. and Guglielmo Weber. 2010. "Consumption and Saving: Models of Intertemporal Allocation and Their Implications for Public Policy." *Journal of Economic Literature* 48:693–751.
- Bajari, Patrick, Phoebe Chan, Dirk Krueger, and Daniel Miller. 2012. "A Dynamic Model of Housing Demand: Evidence and Policy Implications." *International Economic Review* 54:409–442.
- Baker, Scott R. 2013. "Debt and the consumption response to household income shocks." Unpublished manuscript, Stanford Univ., Stanford, CA.
- Banks, James, Richard Blundell, and Agar Brugiavini. 2001. "Risk Pooling, Precautionary Saving and Consumption Growth." *Review of Economic Studies* 68 (4):757–779.
- Berger, David, Veronica Guerrieri, Guido Lorenzino, and Joseph Vavra. 2016. "House Prices and Consumer Spending." Mimeo.
- Bernanke, Ben S, Mark Gertler, and Simon Gilchrist. 1999. "The financial accelerator in a quantitative business cycle framework." *Handbook of Macroeconomics* 1:1341–1393.
- Blundell, Richard, Martin Browning, and Costas Meghir. 1994. "Consumer Demand and the Life-Cycle Allocation of Household Expenditures." *Review of Economic Studies* (61):57–80.
- Blundell, Richard, Hamish Low, and Ian Preston. 2013. "Decomposing Changes in Income Risk Using Consumption Data." *Quantitative Economics* 4:1–37.
- Blundell, Richard, Luigi Pistaferri, and Ian Preston. 2008. "Consumption Inequality and Partial Insurance." *American Economic Review* 98 (5):1887–1921.
- Blundell, Richard and Ian Preston. 1998. "Consumption inequality and income uncertainty." *Quarterly Journal of Economics* 113 (2):603–640.
- Browning, Martin, Angus Deaton, and Margaret Irish. 1985. "A Profitable Approach to Labour Supply and Commodity Demands over the Life-Cycle." *Econometrica* 53 (3):503–544.
- Bunn, Philip and May Rostom. 2014. "Household debt and spending." *Bank of England Quarterly Bulletin* 54 (3):304–315.
- Campbell, John Y and Joao F Cocco. 2007. "How do house prices affect consumption? Evidence from micro data." *Journal of Monetary Economics* 54 (3):591–621.

- Carroll, Christopher D. 1992. “The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence.” *Brookings Papers on Economic Activity* 1992(2):61–156.
- . 1997. “Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis.” *Quarterly Journal of Economics* 112 (1):1–55.
- Carroll, Christopher D. and Andrew A. Samwick. 1997. “The Nature of Precautionary Wealth.” *Journal of Monetary Economics* 40(1):41–71.
- Carroll, Christopher D. and Lawrence H. Summers. 1991. “Primary Surpluses and Sustainable Debt Levels in Emerging Market Countries.” *In National Saving and Economic Performance*, edited by B. Douglas Bernheim and John B. Shoven (Chicago: Univ. Chicago Press (for NBER)).
- Cloyne, James S and Paolo Surico. 2017. “Household debt and the dynamic effects of income tax changes.” *Review of Economic Studies* 84 (1):45–81.
- Cocco, João F., Francisco J. Gomes, and Pascal J. Maenhout. 2005. “Consumption and Portfolio Choice over the Life Cycle.” *Review of Financial Studies* 18(2):491–533.
- Deaton, Angus. 1985a. “Panel Data from the Time Series of Cross-Sections.” *Journal of Econometrics* 30:109–126.
- . 1985b. “Panel Data from the Time Series of Cross-Sections.” *Journal of Econometrics* 30 (1):109–126.
- . 1991. “Saving and Liquidity Constraints.” *Econometrica* 59(5):1221–1248.
- Dynan, Karen. 2012. “Is a Household Debt Overhang Holding Back Consumption?” *Brookings Papers on Economic Activity* 2012 (1):299–362.
- Eggertsson, Gauti B and Paul Krugman. 2012. “Debt, Deleveraging, and the Liquidity trap: A Fisher-Minsky-Koo Approach.” *Quarterly Journal of Economics* 127 (3):1469–1513.
- Fisher, Irving. 1933. “The Debt-Deflation Theory of Great Depressions.” *Econometrica* 1 (4):337–357.
- Gourinchas, Pierre-Olivier and Jonathan A. Parker. 2002. “Consumption over the Life Cycle.” *Econometrica* 70(1):47–89.
- Guerrieri, Veronica and Guido Lorenzoni. 2011. “Credit crises, precautionary savings, and the liquidity trap.” Tech. rep., National Bureau of Economic Research.

- Hall, Robert E. 1978. “Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence.” *Journal of Political Economy* 96:971–987.
- Hurd, Michael D. 1989. “Mortality Risk and Bequest.” *Econometrica* 57:779–813.
- Hurst, Erik and Frank Stafford. 2004. “Home is Where the Equity Is: Liquidity Constraints, Refinancing and Consumption.” *Journal of Money, Credit and Banking* 36 (6):985–1014.
- Iacoviello, Matteo. 2005. “House Prices, Borrowing Constraints, and Monetary Policy in The Business Cycle.” *American Economic Review* 95 (3):739–764.
- Jones, Callum, Thomas Philippon, and Virgiliu Midrigan. 2018. “Household leverage and the recession.” Tech. Rep. 16965, National Bureau of Economic Research.
- Jordà, Òscar, Moritz HP Schularick, and Alan M Taylor. 2013. “Sovereigns versus Banks: Credit, Crises, and Consequences.” Tech. rep., National Bureau of Economic Research.
- Judd, Kenneth. 1998. “Numerical Methods in Economics.” *The MIT Press* Cambridge, Massachusetts.
- Kaplan, Greg, Kurt Mitman, and Giovanni L Violante. 2017. “The housing boom and bust: Model meets evidence.” Tech. rep., National Bureau of Economic Research.
- King, Mervyn. 1994. “Debt deflation: Theory and evidence.” *European Economic Review* 38 (3):419–445.
- Kovacs, Agnes. 2016. “Present Bias, Temptation and Commitment Over the Life-cycle: Estimating and Simulating Gul-Pesendorfer Preferences.” *University of Oxford, Department of Economics Discussion Paper Series, No. 796* .
- Mian, Atif, Kamalesh Rao, and Amir Sufi. 2013. “Household Balance Sheets, Consumption, and the Economic Slump*.” *Quarterly Journal of Economics* 128 (4):1687–1726.
- Midrigan, Virgiliu and Thomas Philippon. 2011. “Household Leverage and The Recession.” *NBER Working Paper* (16965).
- Moffitt, Robert. 1993. “Identification and estimation of dynamic models with a time series of repeated cross-sections.” *Journal of Econometrics* 59 (1):99–123.
- Piazzesi, Monika, Martin Schneider, and Selale Tuzel. 2007. “Housing, Consumption and Asset Pricing.” *Journal of Financial Economics* 83(3):531–569.

- Saiz, Albert. 2010. "The Geographic Determinants of Housing Supply." *Quarterly Journal of Economics* 125 (3):1253–1296.
- Skinner, Jonathan. 1996. "Is Housing Wealth a Sideshow?" *Advances in Economics of Aging*, Wise, D. ed. Chicago: University of Chicago Press:241–268.
- Wong, Arlene. 2015. "Population aging and the transmission of monetary policy to consumption." *Unpublished manuscript* .
- Zeldes, Stephen P. 1989. "Consumption and Liquidity Constraints: An Empirical Investigation." *Journal of Political Economy* 97:305–346.