Consumption spending, housing investments and the role of leverage

Thomas F. Crossley* Peter Levell † Hamish Low ‡

May 26, 2018,

Preliminary

Abstract

In this paper we estimate the degree to which leverage amplifies the effects of house price shocks on consumer borrowing and spending. We show that households re-leverage in response to house price shocks by borrowing more as prices increase. We also find strong response of residential investment spending to house price gains that increase with households initial leverage. By contrast we find limited evidence of housing wealth effects on consumption spending. Our results can be rationalised in a framework where households treat leverage as a portfolio choice, choosing leverage to optimise the risk and return on their assets. When house prices rise, leveraged households respond to the reduction in their loan-to-value ratios by borrowing and investing in their homes in order to re-leverage their portfolios.

Keywords: House prices, leverage, consumption

*Department of Economics, University of Essex Wivenhoe Park, Colchester, UK, C04 3SQ (email: tcross@essex.ac.uk)

†Institute for Fiscal Studies, 7 Ridgmount Street, London WC1E 7AE (email: peter_l@ifs.org.uk).

‡Faculty of Economics, Austin Robinson Building, Sidgwick Avenue, Cambridge, CB3 9DD (email: Hamish.Low@econ.cam.ac.uk). The authors would like to thank Jirka Slacalek, participants of the Bundesbank Household Finance Workshop, participants of the 2017 CESifo Venice Summer Institute, and participants of the 2017 NBER Summer Institute for helpful comments. The research was supported by the ESRC-funded Centre for Microeconomic Analysis of Public Policy at the Institute for Fiscal Studies (CPP, reference RES-544-28-5001). Funding from the EU Horizon 2020 ADEMU project is also gratefully acknowledged.

1 Introduction

Two factors make movements in house prices especially important for consumption spending and the wider economy. The first is that it represents a significant proportion of households' wealth.

The second factor is leverage. The vast majority of US and UK households purchase their first home with substantial credit, leading to a large difference between households gross and net asset positions. As a result, leverage can magnify the effects of economic shocks on households' overall balance sheets, with potentially important implications for their willingness to spend, investment decisions, and access to credit over the business cycle.

A large literature has now developed looking at how house price changes affect consumer spending (see for example Aladangady (2017); Cloyne, Huber, Ilzetzki, and Kleven (2017); Mian, Rao, and Sufi (2013); Mian and Sufi (2011)). Others have looked specifically at the reasons and degree to which household leverage affects these responses (Berger, Guerrieri, Lorenzoni, & Vavra, 2017; Disney, Bridges, & Gathergood, 2010; Disney, Gathergood, & Henley, 2010). These papers have tended to find that leverage magnifies the effect of house price changes, and have a proposed a number of channels which might explain this. These are firstly: a *portfolio* or wealth channel (that more leverage households will see a greater change in housing wealth for a given change in house prices); a *collateral* channel (that high leverage is an indicator of credit constraints); and a *debt-overhang* channel (that households with higher leverage might reduce spending so as to meet some target level of leverage). This latter channel is particularly associated with Dynan (2012).

In this paper we emphasise an alternative mechanism through which leverage might affect household spending decisions. The above mechanisms pertain particularly to *consumption* spending, but house price increases may also affect household's desire to invest in housing. In life-cycle models of household portfolios (e.g. Merton (1969)), desired leverage should evolve smoothly over time. House price decreases will therefore tend to leave households over-leveraged relative to this target, inducing a desire to reduce consumption spending and increase savings. This behaviour provides a micro-foundation for 'debt-overhang' effects. We also highlight an important corollary to this effect: in periods where house prices increase, consumers will find themselves *under*-leveraged. This will have the effect of encouraging consumers to borrow and invest in housing stock. Such portfolio adjustments increase leverage and return the desired rate and variance of returns, but they increase the *size* of household balance sheets. Since portfolio shares are themselves nonlinear functions of leverage, we would expect these effects to be greater for households with higher loan-to-value ratios. More leveraged households will experience a larger reduction in the portfolio share of housing for a given price increase and so we predict that they should respond more strongly than other households.

We test this hypothesis empirically in two ways. First, we provide panel-data evidence on the extent of household re-leveraging in response to house price shocks, demonstrating that households tend to borrow against rising home values. Household mortgage debt among existing home-owners, who do not move home, increases by roughly 0.3 percentage points with each 1% increase in home values in the US and by 0.2 percentage points in the UK. In this respect, households increased the size of their balance sheets in the years prior to the financial crisis in a similar way to investment banks (Adrian & Shin, 2010).

Second, we examine how the consumption and residential investment responses of households to house price increases varies according to household leverage. To test our hypothesis that more leveraged households respond to price rises by increasing residential investment spending, we require detailed data on consumer spending broken down into its separate consumption and investment components. However, datasets that do this typically do not include information on household's wealth or leverage. Those surveys that do (such as the Panel Study of Income Dynamics), typically do not yet cover long time periods. We therefore combine information from two surveys. The first of these contains information on households balance sheets but which does not have detailed consumption data. The second is a national budget survey which, in common with many surveys of consumer spending, contains information on households expenditures but not on their wealth. Our approach is to apply of two-sample IV methods (Angrist & Krueger, 1992) using a source of variation in leverage that is common to both samples. The instrument we use is the a measure of credit conditions when households first moved into their current homes (the average loan-to-income ratio on new house purchases). We show using evidence from our first sample that this instrument is conditionally uncorrelated with a number of potential confounding variables including gross house values, financial asset income, and unsecured debt.

Using this approach, we find that leveraged returns are associated with rapid increases in residential investment spending by existing homeowners. We estimate an elasticity of residential investment spending to house price increases of 0.75 among outright owners. The estimated elasticity is twice as large for those with a loan-to-value ratio of 50% (and scales linearly with each further doubling of households' debt to equity ratios). Responses for total spending and non-durable spending are by contrast proportionally much smaller and not statistically different from zero.

Much of the previous literature looking at questions relating to leverage and household spending has attributed larger spending responses among more leverage households to the relaxation of credit constraints. The typical empirical approach to the question of how and why leverage affects household responses to housing wealth shocks, is to run a regression of the form

$$\Delta c_{it} = \beta_0 + \beta_1 \Delta p_{it} + \beta_2 \Delta p_{it} \times f(L_{it-1}) + \epsilon \tag{1}$$

Here c_{it} is some *measure* or indicator of consumption spending, saving or borrowing, p_{it} is the individual's house price and $L_{i,t-1}$ is the individual's lagged leverage (the loan to value - LTV - ratio on the individual's home). The function $f(.)$ is sometimes a piece-wise linear spline allowing the effects of leverage on house price changes to vary in a nonlinear way. c_{it} is often only an indicator of total spending (or related measure such as borrowing) as comprehensive consumption measures are rarely included in surveys that contain panel data on wealth. Moreover, it often lumps spending on nondurable consumption together with spending on durables and residential investment (thus providing estimates of the marginal propensity to spend out of increased housing wealth rather than to consume).

The most common finding is that β_2 is positive, but that the amplifying effects of leverage tend to be concentrated among those households who have the very highest leverage levels. The finding that house price effects are highly nonlinear in leverage is typically interpreted as evidence of credit constraints and the primary importance of the collateral channel in explaining house price effects. For example, Disney, Gathergood, and Henley (2010) find that savings respond very little on average to house price shocks but that the response is five times greater for households emerging from a situation of negative equity than for households with initially positive equity values. Cooper (2013) finds that the spending of households with high debt service ratios, low liquid wealth and high expected future income were the most sensitive to housing wealth changes. Since all these characteristics are indicative of credit constraints, he argues that this is the most important channel at work. Looking at the UK for the years 1995-2005, Disney, Bridges, and Gathergood (2010) also find that households' propensity to remortgage in the face of price rises does not increase with household's LTV ratios but is higher for those with high LTV *and* high unsecured debt. In a related paper that using the US Panel Study of Income Dynamics, (PSID) Disney and Gathergood (2011) again look at changes in household indebtedness by initial LTV ratios (again allowing these to be interacted with levels of unsecured debt). As before, households with high leverage - particularly those with high leverage and high unsecured debt - were seen to be much more responsive. Mian and Sufi (2011) examine the link between growth in debt and regional variation in home prices in the United States over the period 2002-2006. They instrument for changes in house prices using estimates of the elasticity of housing supply in different regions (reasoning that regions with more elastic supply should see smaller price increases for a given demand shock). Their estimates suggest very strong effects of house price growth on borrowing. Each 1% gain in house prices is estimated to increase borrowing by 0.52 percent. In addition, they find that consumers with credit scores one standard deviation above the 1997 mean have estimated responses that are twice as large as those with estimated responses one standard deviation below the mean (with an elasticity of 0.75 compared to 0.35). They find no evidence of house price effects for those in the top quartile of the credit score distribution. They also find that responses are in general larger for younger households and find similar differences between consumers with high and low credit utilisation rates. In line with the rest of the literature, they argue that these differences most likely reflect the importance of credit constraints.

In contrast to the previous literature, we argue that, while some households undoubtedly face credit constraints, they are unlikely to drive the effects we document. Our instrumental variable strategy captures differences in leverage across households that result from past decisions (the timing of moves). This is unlikely to pick up differences in the propensity to be credit constrained, and we find no relationship between our instrument and measures of unsecured debt or other forms of financial wealth. In addition, we find no evidence that leverage increases non-durable or durable consumption responses to house prices, but rather that effects are concentrated among residential investment spending. Such behaviour would be difficult to explain through credit constraints alone. Finally, we do not see very different effects in estimates for a younger sub-sample of households than we do in our main sample. As a result, we see our estimates as providing evidence for our alternative portfolio rebalancing channel.

The remainder of this paper is structured as follows. In Section 2 we set-out a theoretical framework explain household's choice of leverage, motivations for re-leveraging and spending decisions. Section 3 describes the data we will use. In Section 4 we provide evidence that households do indeed re-leverage by increasing borrowing when house prices rise. In Section 5 we set out our empirical approach to identify the role leverage plays in household spending decisions, and present results on the impact of house prices on different components of consumption spending. Section 6 concludes.

2 Framework

2.1 A Simple Model

Consider a household with two assets available to it. Housing is a risky asset (h_t) with price p_t. The return on housing is:

$$r_t^* = \frac{p_t}{p_{t-1}} - 1. \tag{2}$$

There is also a risk-free asset (a bond) denoted b_t with price 1 and interest rate r. The household can short the bond (that is, take a loan) and there are no credit constraints. The household cannot short housing.

Leverage (the loan-to-value ratio) is:

$$L_t = \frac{\text{debt}}{\text{house value}} = \frac{\text{debt}}{\text{gross housing wealth}} = \frac{-b_t}{p_t h_t} \tag{3}$$

and the portfolio share of housing is:

$$\omega_t = \frac{\text{gross housing wealth}}{\text{net wealth}} = \frac{p_t h_t}{p_t h_t + b_t} = \frac{1}{(1 - L_t)} \tag{4}$$

Leverage $0 < L_t < 1$ implies $\omega_t > 1$. For example, a household with a 95% "mortgage" ($L = 0.95$) has a housing portfolio share of $\omega_t = 20$, while for outright owners $\omega_t = 1$.

Leverage magnifies risk and return. Denote the household's net wealth by x_t, consumption by c_t, and labour income by y_t . Then the intertemporal budget constrant is:

$$x_t = (1 + r + \omega_{t-1} (r_t^* - r)) * (x_{t-1} - c_{t-1}) + y_t \tag{5}$$

$$x_t = \left(1 + r + \frac{1}{1 - L_{t-1}} (r_t^* - r)\right) * (x_{t-1} - c_{t-1}) + y_t \tag{6}$$

Leveraged households have a greater wealth gain for a given house price shock, and these effects are highly nonlinear in leverage. Note though, that the variance of portfolio returns are also higher for leveraged households.

To go further we need a model of consumption and portfolio choice. We begin with a very simple one. For the moment, assume the household has no labour income ($y_t = 0$) and there are no adjustment costs associated with housing. Further assume that the household has CRRA preferences over consumption and that housing is just an investment good (does not yield a flow of utility). This is essentially the life-cyle portfolio choice model of Merton (1969)), and the policy functions are well known. There is a linear consumption function:

$$c_t = \alpha_t x_t \tag{7}$$

and a constant target portfolio share for the risky asset:

$$\omega_t = \omega^* \tag{8}$$

In the Merton model, the portfolio share of the risky asset depends only on moments of the return distribution. As leverage is just a transformation of the housing portfolio share, this implies there is a constant target leverage that delivers the household's desired combination of risk and return.

Now consider how this household responds to a positive house price shock:

$$x_t - E[x_t] = \omega^* (r_t^* - E[r^*]) * (x_{t-1} - c_{t-1}) \tag{9}$$

$$x_t - E[x_t] = (p_t - E[p_t]) * h_{t-1}. \tag{10}$$

The shock is partly consumed:

$$c_t - E[c_t] = \alpha_t (x_t - E[x_t]) \tag{11}$$

and partly saved (s_t) according to the consumption function:

$$s_t - E[s_t] = (1 - \alpha_t)(x_t - E[x_t]). \tag{12}$$

Critically, the portfolio choice rule implies that this additional saving is leveraged by ω^*:

$$p_t * h_t - E[p_t * h_t] = \omega^* \left(1 - \alpha_t\right) \left(x_t - E[x_t]\right). \tag{13}$$

Using:

$$x_t - E[x_t] = \left(p_t - E\left[p_t\right]\right) * h_{t-1} \tag{14}$$

this implies extra active investment in housing of:

$$\left(p_t * h_t - E[p_t * h_t]\right) - \left(p_t * h_{t-1} - E[p_t] * h_{t-1}\right) = \left(\omega * \left(1 - \alpha_t\right) - 1\right) \left(p_t - E[p_t]\right) h_{t-1} \tag{15}$$

The first term on the left hand side of equation (15) is the desired extra gross housing wealth. The second term is the additional housing wealth that comes mechanically from the unexpected price increase. The difference between the two is the additional active investment in housing (funded by debt) to return the housing portfolio share to w^*. Thus, if there is an unexpected house price increase, a leveraged household will invest more in housing and borrow to do so (even if the household believes that housing returns are iid.) Conversely, an unexpected house price fall increasing the leverage of the portfolio and the household will want to sell housing and retire debt to return to ω^*.

For example, suppose that the household owns a £600,000 house with $\alpha = 0.05$ and $\omega = 3$ (so that the household has 33% equity in the home.) If the house value unexpectedly goes up by 5% (£30,000), the consumption function implies that net wealth increases by £28,500 and the constant portfolio rule implies that the household then desires gross housing wealth of £685,500 . As the house value is now £630,000, the households makes new investment in housing of £55,500, financed by new debt. Note that the extra investment spending (£55,500) is much larger than the extra consumption spending (£1,500).

This model is *very* stark, and the mechanism described will be moderated in a number of ways. As one example, if housing provides a flow of utility, this is like a dividend, but if there is diminishing marginal utility from housing (and extra housing cannot be rented efficiently), the total return to housing falls with housing wealth held, and this will temper some of the releveraging motive described above.

Nevertheless, equations (12) and (13) suggest that the re-leveraging mechanism we describe will operate so long as the policy functions for consumption and for the risky asset portfolio share are sufficiently flat in net wealth. A gently sloped consumption function implies that a significant fraction of a wealth shock is saved. A fairly flat portfolio rule implies that the household will not want to change its portfolio shares dramatically in response to a wealth shock. A wide range of life-cycle consumption and portfolio-choice models share these features.

We therefore do two things in the remainder of the paper. First, in in the next subsection we illustrate by numerically solving a much richer model with idiosyncratic labour income risk, utility flows from housing, and transaction costs in adjusting housing. Then, in the empirical part of the paper, we investigate the empirical relevance of this mechanism.

2.2 A More Realistic Model

To Come

3 Data

To investigate re-leveraging and the relationship between leverage and household spending, we draw on three datasets.

The first is the Living Costs and Food Survey and it's previous incarnations the Expenditure and Food Survey and Family Expenditure Survey (which we shall refer to collectively as the LCFS). The LCFS is a comprehensive, long-running survey of consumer expenditures involving between 5,000-8,000 households per year. Households are asked to record high-frequency expenditures in spending diaries over a two week period. Recall interviews are used to obtain spending on information on big ticket items (such as holidays or large durables) as well as standing costs on items such energy and water, internet bills and magazine subscriptions. The survey also collects information on incomes, demographic characteristics and, since 1992, on the value of households' mortgages (but not on other aspects of household balance sheets such as home values).

The second dataset we use is the British Household Panel Survey and its successor Understanding Society (both of which we shall refer to as the BHPS) (University of Essex. Institute for Social and Economic Research (2010); University of Essex. Institute for Social and Economic Research. (2016)). The BHPS is available in 18 waves from 1991 to 2008. Understanding Society began in 2009 and incorporated the original BHPS sample members from 2010 onwards. Both surveys include limited information on household spending on food and drink as well as self-reported house values. The BHPS contains data on mortgage values from 1993 onwards, while Understanding Society dropped these variables in its second wave in 2010. In the remaining years, we continue to observe whether households own their homes outright, and details on the length and type of their mortgage if they have one. We use these along with past information on mortgages values to impute mortgages in years following 2010. Loan to value ratios are calculated by dividing the value of mortgages by the (self-reported) value of homes.

The need to use two UK surveys comes from the fact that consumption spending is observed in the LCFS but leverage is not. At the same time the BHPS includes information on leverage but not on consumer spending. Hence our need to use two-sample methods that combine the information contained in both datasets, as we describe below.

The third dataset we use is the Panel Study of Income Dynamics (PSID). The PSID is a US-based panel of households that includes information on home ownership, household balance sheets, income and spending decisions. Since 1997, the survey has been biennial. The PSID has included questions on the value of households home equity and mortgages on an annual basis from 1999 onwards. Prior to 1999, these were only asked every 5 years. In terms of spending data, the survey only consistently included spending on food and rental payments until 1999. In that year, this was extended to cover other non-durable expenditures (including health, utilities, education and childcare). Other expenditures such as clothing and entertainment were added in 2005. Since 2001, households have were asked whether they have undertaken home improvements worth $10,000 or more since January of the year two years prior to the interview. If they answer in the affirmative, they are then asked to give the exact amount spent.[1]

For house prices we use regional/state-level data on the prices of transacted houses published by the Office for National Statistics (for the UK) and the Federal Housing Finance Agency (for the US).

In all of what follows, we drop households where the head is aged under 25 or over 65. To avoid problems of measurement error when estimating our first stage, we also drop households who have a lagged housing portfolio share in the top 1% of the distribution and those who have negative equity. We also drop households resident in Northern Ireland from both the BHPS and the LCFS samples as these were only introduced into the BHPS sample in later years. Finally, to avoid our results being partially driven by households moving at times of high or low house price growth, we drop households who have lived in their home for less than one year.

Table 1 provides some descriptive statistics for our three samples. We report these for 1993-2013 in the UK data, and for 2005-2013 for the PSID (i.e the years when the most comprehensive spending data was available). The proportion of those owning their own homes and the average tenure among homeowners are similar across the two UK surveys, at around 70% of households. Ownership rates are somewhat lower in the PSID at around 55%. Focussing on home owners, the average loan-to-value ratio in our BHPS sample is 0.34, while US households tend to be more levered with an average loan-to-value ratio of 0.54. As we see below this is partly accounted for by the fact the age gradient in leverage is less steep in the United States, meaning that older households in the US tend to have higher loan-to-value ratios than their UK counterparts.

Figures from the LCFS show that non-durables is the largest component of expenditure (accounting for 76%). Residential investment spending, which includes extensions, renovations, household repairs, large furniture, carpets, and large household appliances accounts for roughly 7% of total spending. The remainder is accounted for by spending on non-residential durables.

[1] We annualise this figure using individuals' month of interview to determine the exact length of the period covered by this question.

The spending questions included PSID are as detailed as in the LCFS, and so we are forced to define categories differently for the US. We measure residential investment in the PSID as the sum of responses to the question "how much did your family spend altogether on household furnishings and equipment, including household textiles, furniture, floor coverings, major appliances, small appliances and miscellaneous housewares?" and responses to questions regarding home improvement spending (which are censored from below at $10,000). Since we are unable to exclude spending on small furnishings and smaller electrical appliances from this value, this definition is somewhat broader than the one used in the UK. As measured, it accounts for 6.2% of total spending. Non-residential durable spending in the PSID is essentially restricted to cars. Relative to our definition for the LCFS, it therefore excludes audio-visual equipment, as well appliances such as vacuum cleaners and microwaves (which may be included as durable household furnishings). This category accounts for 9.9% of expenditures. The remaining 83.9% of measured spending (including clothing, utilities, entertainment, vacations, motor fuel, healthcare and child care) is classified as going on non-durables. In all of these categories, non-responses to individual questions are treated as implying zero expenditures.

Table 1: Descriptive statistics, BHPS and LCFS and PSID

	BHPS	LCFS	PSID
	(1993-2013)	(1993-2013)	(2005-2013)
Age	44.5	44.2	42.7
% Own home	68.7%	70.3%	54.5%
Homeowners			
Years at address	11.0	10.2	11.1
LTV ratio	0.34	-	0.46
ω_t (housing share)	3.04	-	4.06
Total spend ($ ann.)	-	43,764	65,518
Non-durable	-	33,430	54,977
Durable	-	6,068	6,495
Residential inv.	-	4,266	4,407
% Res inv. > 0	-	79.6%	73.1 %

Note: UK data is for the period 1993-2013. US data is for the period 2005-2013 when more comprehensive spending measures are available in the PSID. See text for details of what is included in each spending category.

4 Evidence for re-leveraging

We now turn to describing how the leverage of households in the UK and US has evolved over time and, in particular, how it has responded to increases in house prices in the two countries. We begin with aggregates trends in leverage and house prices, before moving on to household-level panel regressions describing within household re-leveraging and de-leveraging behaviour.

Figure 2 shows how leverage evolves over time across four different 10-year birth cohorts (born between the years 1940 and 1970). In both the UK and the US there is a steady and reasonably smooth decline in leverage by age. In the UK there are pronounced differences in leverage between cohorts at younger ages. However, the different cohorts largely converge to similar leverage by around age 45. As we discuss further below, the differences in initial leverage across UK cohorts are likely to be explained by the differing credit conditions and house prices the different cohorts were exposed to at the point they became home-owners. This is a source of variation in leverage which we exploit to identify the role of leverage in spending decisions in what follows.

In the US, there is much less evidence of cohort effects, and the decline in leverage by age is much less steep than in the UK.

Figure 1: LTV ratios by age and cohort

(a) UK (1993-2013) (b) US (1999-2013)

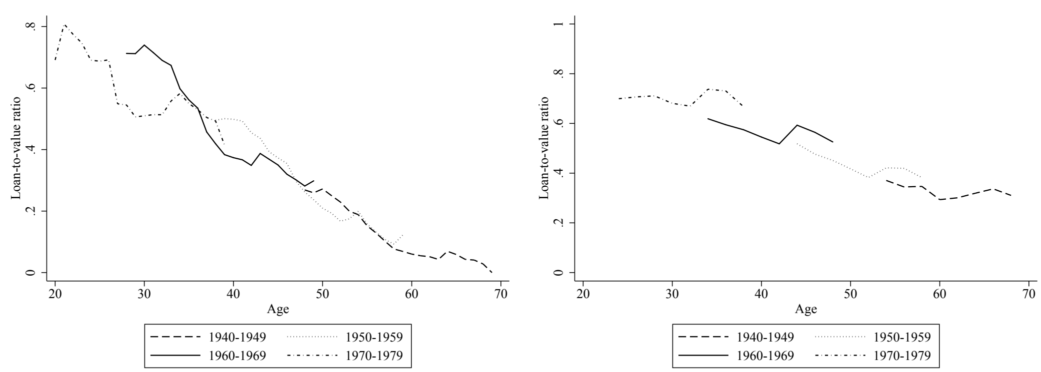

Source: Authors' calculations using BHPS/Understanding Society and PSID

Next we consider how leverage varies over the house price cycle. The two panels of Figure 2 shows how average real house price growth varied from year to year in the UK and US, as well as changes in households' average loan-to-value ratios among both younger (aged 25-45) and older (45-65) households in the two countries.

In the US, loan-to-value ratios among both younger and older households remained strikingly stable throughout the period of house price growth that continued until 2006. This peaked at a national rate of 7.6% in 2005, when loan-to-value ratios were essentially unchanged from the previous year (at around 60% for those aged 25-45 and 40% for those aged 46-65). When real house prices started to decline from 2007 onwards however, loan-to-value ratios rose rapidly. House prices fell by 2% in 2007 and between 7-8% in each of the years from 2008-2011. Over this period, the average LTV among younger households increased from 62-71%, while for older households it increased from 37-44%.

Such declines did not occur in the UK, where the house price slump was modest relative to both previous UK house price slumps and to the declines observed in the US (Bénétrix, Eichengreen, & O'Rourke, 2012). We plot UK house price and loan-to-value data from 1993-2013. For most of this period, UK house prices were increasing, with annual falls only observed in 1994-1995 and 2007-2009. In the period in between these years, house prices grew rapidly. Annual price increases peaked in 2003 at a rate of almost 20%. There is more evidence of a fall in average loan-to-value ratios as house prices rose in the UK than in the US. Loan-to-value ratios fell somewhat in the period of greatest house price growth among the under 45s, falling from 62% in 1995 to 43% in 2004 before climbing again as house price growth moderated (the over 45s saw smaller changes in their average leverage). However, as we now discuss, there is evidence that UK households were also engaging in re-leveraging behaviour over this period, even if the scale was not as great as it was in the US.

Figure 2: LTV ratios and house price growth rates

(a) UK (1993-2013) (b) US (1999-2013)

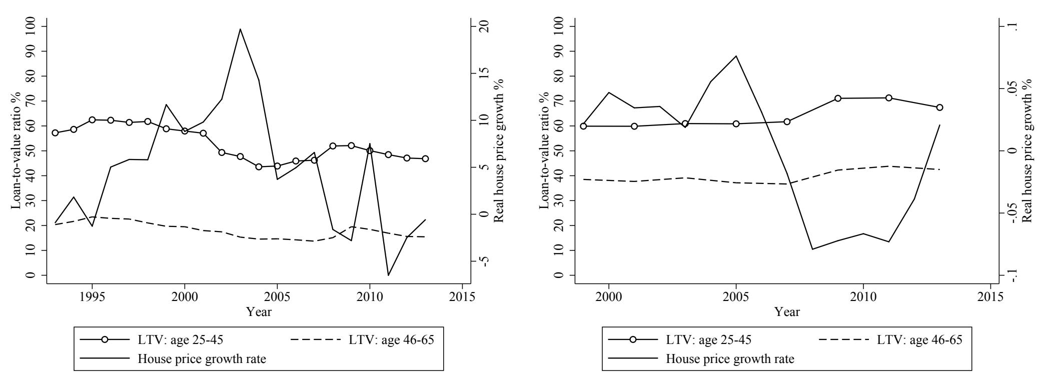

Notes: House prices for the UK are national averages taken from the Office for National Statistics HPI deflated using the UK CPI. House prices in the US are national averages taken from the Federal Housing Finance Agency and are deflated with the US CPI. Loan-to-value ratios are taken calculated for the UK using data from BHPS and Understanding Society and for the US using the PSID.

Panel (a) in Figure 3 shows that younger households in the UK did not passively allow leverage to fall as prices rose. Rather they responded by increasing the amount of new borrowing. The proportion of home-owning households aged 25-45 observed taking out additional mortgage debt in the British Household Panel Survey increased to exceed 10% in the period of most rapid house price growth. Panel (b) of Figure 3 shows that this period of new borrowing also coincided with growth in the proportion of households with positive spending on housing extensions. This activity was also focused on younger households. As we saw in Figure 2, older households responded much less to these developments. The importance of extensions and other home improvements as a reason for new borrowing is confirmed when we consider the uses for which households report taking out additional mortgage debt. Households in the BHPS are asked whether new loans were used for extensions, home improvements, car purchases or other household spending (households may give more than one answer). We group these responses into “residential investment’ and “consumption” and plot the proportions reporting each motive for household heads aged 25-45, and 46-65 in panels a and b of in Figure 4. Younger households are roughly four times more likely to report taking out a loan for residential investment than for consumption spending, and this motive for re-leveraging also appears to be the most pro-cyclical. Housing investment is also a more important motive for older households to take out additional mortgage debt.

Figure 3: New mortgage loans and housing extensions by homeowners in the UK, 1993-2013

(a) Proportion taking out additional mortgages

(b) Proportion with positive spending on housing extensions

Source: British Household Panel Survey/Understanding Society and Living Costs and Food Survey

Figure 4: Purpose of new mortgage loans, 1993-2013

(a) Aged 25-45 (b) Aged 46-65

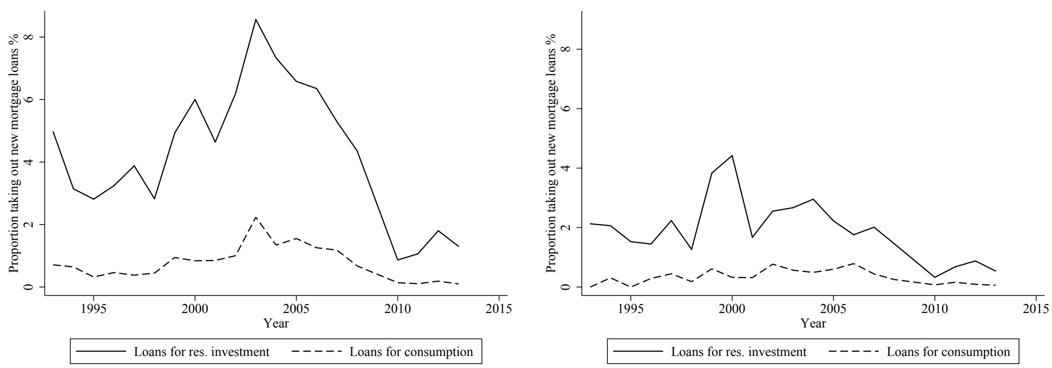

Source: British Household Panel Survey/Understanding Society

Figure 5 shows trends the proportion of home-owners engaged in new mortgage borrowing alongside changes in the proportion of households undertaking home improvement spending for the PSID (both since the previous wave - i.e in the previous two years). As in the UK, younger households in particular were more likely to take out new mortgage loans in periods of high house price growth. There is also evidence of increases in home improvement spending when house prices growth was highest in between the 2003 and 2007 waves of the survey. Unlike in the UK, however, this is not clearly greater for younger households. While the PSID does not include questions on the motives for additional mortgage borrowing as in the BHPS, previous studies have pointed to home improvement spending as a key motive for equity withdrawal (Cooper, 2010).

Figure 5: New mortgage loans and housing extensions by homeowners since previous PSID wave, 1999-2013

(a) Proportion taking out additional mortgages

(b) Proportion with positive spending on housing extensions

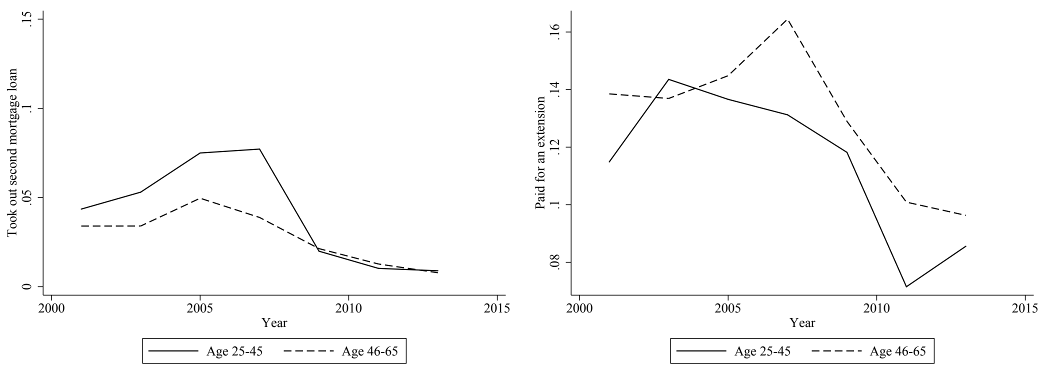

Source: Panel Study of Income Dynamics

Taking these trends together, Figure 6 shows trends in average household debt and equity (normalised by income) in the two countries. This illustrates the increase in the size of household balance sheets over the period of the house price boom, with increases in home equity (house values less mortgages) accompanied by increases in household debt.[2] The increases in home equity were much larger in the UK both relative to income and to average mortgage debt than they were in the US. In 2013, the average home equity to income ratio in the UK was five, compared to three in 1993. In the US, the average home-equity to income ratio was 1.5 in 2013 (compared to around 1 in 1999). We also note how debt burdens in the US decreased only slightly as home equity fell in the years following the financial crisis. This is suggestive of an asymmetry in the relative ease of re-leveraging in response to house price increases and de-leveraging in response to house price falls.

[2]Here home equity is calculated using data on the value of the primary residence.

Figure 6: Housing equity and debt (normalised by income)

(a) UK (b) USA

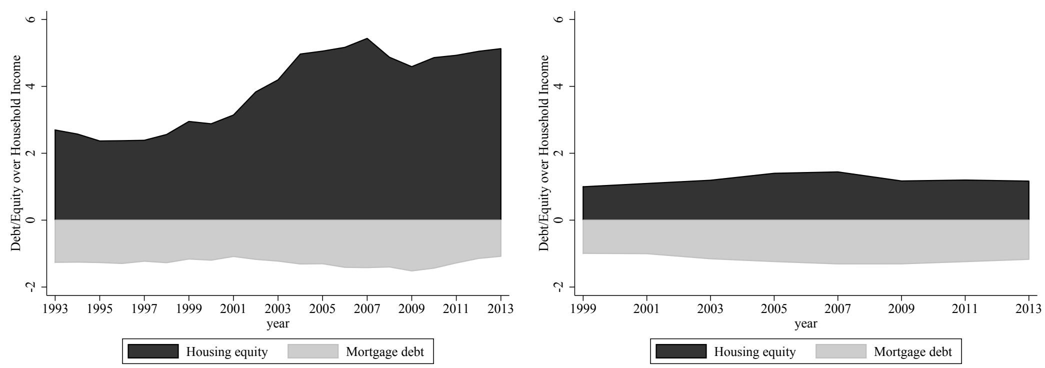

Source: Authors' calculations using the BHPS and PSID

To more formally examine the extent to which households adjusted their leverage as prices rosem, we report in Table 2 results from regressing changes in log mortgage debt on changes in households' log self-reported home values. That is, we run the regression equation

$$\Delta \log Debt_{it} = \delta \Delta \log HValue_{it} + \varepsilon_{it} \tag{16}$$

on a sample of home-owners who are never observed with an LTV of zero. This allows us to directly examine within household response to changes in house values over time. We report results from both the BHPS and the PSID and since the PSID has been biennial since 1999, in both surveys we consider changes over the previous two years.

If mortgage debt did not adjust as house prices increased, leaving LTV ratios to fall passively with house prices, then we would expect the coefficient on log house values to equal 0.

Panel (a) of Table 2 presents these results for the US. Column (1) shows the results for a simple OLS regression, initially focussing on the sub-sample of households who do not move from one wave to the next. The results indicate that each 10% increase in house prices appears associated with a 3.4% increase in mortgage debt. Measurement error in self-reported house values has the potential to attenuate this coefficient towards zero. Column (2) therefore shows results when we instrument the change in house prices with state-level house price growth. The results are very similar, suggesting that measurement error is not a significant problem. The results also do not change much when we include a control for age to account for life-cycle changes in leverage in Column (3). As a result we report only OLS results for the subsequent columns. In Columns (4) and (5) we look for evidence in asymmetries of responses when households are re-leveraging versus de-leveraging by splitting the sample according to whether real regional house prices rose or fell relative to the previous wave. The coefficient on house price changes when prices are rising is significantly greater than at times when house prices are falling. Our results imply that a 10% increase in US house values was associated with a 3.8% increse in mortgage debt, while a 10% fall in house prices is associated with only a 2.5% reduction in mortgage debt. In the remaining three columns we consider how these results change when we include households who move address from one wave to the next. We estimate that debt responses to house price changes are nearly twice as large when movers are included. This holds true when we separately consider responses to house price falls and house price increases in Columns (5) and (6), indicating that up-sizing and down-sizing are important mechanisms by which households adjust their leverage.

Panel (b) shows equivalent responses estimated from the BHPS from 1993-2009. Here we do not include data from Understanding Society, so as to exclude the changes in mortgage debt which we have imputed from 2011 onwards. In the UK, debt responses to house price gains are not as large as they are in the US but remain substantial. The results in Column (1) indicate that each 10% increase in house prices is associated with a 2.0% increase in mortgage debt. As in the PSID, these results do not change much when we instrument for changes in house value (Column (2)) or include an age control (Column (3)). Unlike in the US, we do not find clear evidence of asymmetric responses. The estimated effect of changes in house values on mortgage debt is smaller in response to negative changes than positive ones, but it is not possible to draw clear conclusions for this as, owing to the limited number and duration of episodes of falling UK house prices over the period we consider, the coefficient estimated in Column (4) is very imprecise. When we include movers, the estimated responses of mortgage debt to house price changes rises as it does in the US, though not quite to the same levels. In the sample which includes movers, a 10% inrease in house prices is associated with a 4% increase in debt compared to 6% in the US.

Table 2: Changes in mortgage debt and changes in house prices over two-year periods

$\Delta logDebt$	OLS	IV	IV	OLS	OLS	OLS	OLS	OLS
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel a: US 1999-2013								
$\Delta \log HValue$	0.344***	0.326***	0.323***	0.253***	0.380***	0.602***	0.539***	0.647***
	(0.058)	(0.038)	(0.038)	(0.042)	(0.097)	(0.052)	(0.038)	(0.089)
R^2	0.031	0.031	0.031	0.012	0.039	0.118	0.080	0.135
N	10,811	10,777	10,777	5,088	5,723	11,895	5,489	6,406
Clusters	3,358	3,350	3,350	2,430	2,620	3,477	2,534	2753
Panel b: UK 1993-2009								
$\Delta \log HValue$	0.200***	0.238***	0.238***	0.239	0.196***	0.413***	0.571***	0.411***
	(0.024)	(0.037)	(0.037)	(0.154)	(0.024)	(0.035)	(0.179)	(0.036)
R^2	0.012	0.012	0.016	0.007	0.012	0.052	0.055	0.052
N	12,928	12,860	12,860	541	12,387	14,131	586	13,545
Clusters	2,713	2,708	2,708	398	2,670	2,763	422	2,719
Controls								
Age			X					
Restrictions								
Non-movers	X	X	X	X	X			
House price growth< 0				X			X	
House price growth> 0					X			X

Note: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses. Standard errors clustered at the individual level. The log change in house values is instrumented with the log change in nominal house prices within each region.

The results show evidence of re-leveraging behaviour in the face of price increases, as we would expect in a model where consumers treated leverage as a portfolio choice. We also find evidence that de-leveraging responses in the face of price falls are slower than re-leveraging in response to house price gains.

Before proceeding, it is worth considering a few of the reasons why the re-leveraging responses we estimates in the US appear so much larger than they do in the UK. A number of differences might explain this. The first thing to notes is that house price increases tended to be much greater in the UK, meaning that larger adjustments in mortgage debt would have been necessary in order to maintain constant loan-to-value ratios. A second set of reasons concern the differing institutions across the two countries. In some US states, mortgage loans are non-recourse, meaning that lenders cannot pursue debts that are not covered through sales of foreclosed properties. This may make borrowers more comfortable with the risk of negative equity, since the costs of default in this situation are smaller. Ghent and Kudlyak (2011) for instance find that the monthly probability of default for borrowers in a state of negative equity is 32% higher in states where the is no threat of recourse. In the UK, mortgage loans are recourse loans. Secondly, in the US, mortgage interest is tax deductible, creating more of an incentive to both pay-off mortgages less quickly and to increase mortgage debt when prices rise. This was only true to a very limited extent in the UK during the period we consider. Mortgages were only tax deductible up to a fixed nominal cap of 30,000 (far below the average house price). From 1994 the size of the deduction was limited to 20% falling to 10% in 1995. The tax deduction was eliminated in 1999. Using variability in the importance of tax deductibility in the UK over time, Henderschott, Pryce, and White (2003) find that it can have substantial effects on households' initial loan-to-value ratios. It may also have similar effects on incentives to re-lever, and so partly explain some of the differences in behaviour we observe between the US and the UK.

5 Spending regressions

To test the specific hypothesis that more leveraged households will disproportionately increase housing investment in response to house price increases, we now turn to estimating equation (**??**).

One concern about directly estimating this equation is that leverage is potentially endogenous. The conventional approach to estimating leverage effects is to use individuals' lagged leverage in equations such as (1). As we have seen however, lagged leverage is itself a choice variable that will depend on households expectations of the risk and returns associated with housing investments at the time when consumption and other spending decisions are made. As we document in what follows, lagged leverage is also correlated with gross house values and income from non-housing assets. In order for our empirical application to identify the effects of independently varying leverage, these other variables ought to be held constant. For this reason we instrument leverage using housing market conditions at the time individuals moved into their homes.

A second issue concerns data availability. Long-running surveys that contain balance sheet data on wealth and leverage rarely contain comprehensive consumption measures. A panel survey is required in order to know the consumer's lagged leverage position ω_{it-1}. If a suitable panel were available, it would also allow us to represent $g(\lambda_{t-1})$ with a household fixed effect (or similar) in place of a set of proxies.

Previous studies have addressed this problem by either using available proxies for consumption (such as borrowing, (Mian & Sufi, 2011)), subsets of consumption that are observed (e.g. (Lehnert, 2004)) or measures backed out from the consumer's budget constraint (using the difference between observed income and wealth changes, as in Cooper (2013)). Each of these approaches has drawbacks. Changes in particular categories of consumption, or variables related to consumption need not give the full spending response to shocks. They also do not allow us to investigate how the composition of spending varies as house prices change, which is crucial for allowing us to test the importance of portfolio rebalancing effects. In addition, the use of the budget constraint identity to impute consumption can in general lead to biased estimates of wealth effects in the presence of measurement error (Browning, Crossley, & Winter, 2014). If reported wealth in the previous period is smaller than actual wealth, then leverage as observed by the researcher in that period will be too high and consumption in the current period be too large, biasing estimates upward.

Another option is to use a survey such as the PSID, which does contain both sets of variables in some years. However, the number of waves in which it includes comprehensive consumption data is relatively short (a problem that is also applies to other panel surveys such as the HRS, used by Christelis, Georgarakos, and Japelli (2015) to study questions around leverage). In addition, owing to the fact that the leverage of US households is less sensitive to issues surrounding the time of purchase than it is in the UK, we find that our instruments are much weaker in this setting. For this reason we focus on the UK, combining the LCFS and BHPS using two-sample IV methods (Angrist & Krueger, 1992). This approach allows us to simultaneously impute and instrument for leverage in our (cross-sectional) UK expenditure dataset using balance sheet data taken from the BHPS. Details of our approach are provided in Appendix A.

5.1 Instrument and identification

For our proposed method we require a source of variation in leverage that explains why some households took out larger loans than others that is common to both the BHPS and the LCFS. For this purpose we exploit variation in the average price to income ratios for new loans at the time households moved into their current residences (denoted P/Y_{-T}). This variable is often used as a measure of the cost of credit (*loan* to income ratios for example included in the credit conditions index of (Fernandez-Corugedo & Muellbauer, 2006)). In our case it indicates the cost of borrowing in the years house prices were made, and so the degree to which households would have been able to leverage their housing purchases at the time they moved.

The solid line in Figure 7 (Panel (a)) shows how this instrument varies over time in the UK. There is a gradual upward trend in the price to income ratio suggesting that credit has become looser over time. In 2013, average loans were almost five times greater than the incomes of buyers. This compares to a ratio of 2.5 in 1969. This provides one source of identification. Importantly however, there is also

Figure 7: Credit conditions, 1969-2013

(a) Price-income ratios (b) Advance-income ratios

Source: Office for National Statistics

cyclical variation in this variable, with for example evidence of credit tightening following the 2008 financial crisis. Movements in other measures of credit conditions such as the average deposit on new homes (Figure 7, Panel (b)) show similar patterns.

Our instrument is only available from 1969 onwards, and so in what follows we drop households who moved into their homes before this. This constitutes roughly 0.5% of the total number of observations in our LCFS sample.

5.1.1 Instrument relevance and validity

For our instrument to be considered appropriate it must satisfy two requirements. The first is that the instrument must be relevant (that it is indeed correlated with the endogenous variables it is replacing). Figure 8 shows how our instrument relates to loan-to-value ratios for a given cohort (those born in the 1960s). This is the only ten-year birth cohort that we observe for almost our entire sample period. We plot loan-to-value ratios for households who moved into their homes in three different years: 1989,1996, and 2004. These three years represent peaks and troughs in price-to-income ratios on new housing purchases from Panel (a) in Figure 7. Price to income ratios reached a temporary high of 3.7 in 1989 before falling to a low of 3.2 in 1996. Thereafter they increased to a peak of 5.2 in 2004. As Figure 8 shows, households that moved when price-to-income ratios were relatively high in 1989 tended to have higher leverage than those in the same cohort who moved in in 1996. This is true not only at the point they moved in to their current homes but also long-afterward. Loan-to-value ratios are also persistently higher for those who moved in when credit conditions were even looser in 2004.

This relevance of our instruments can be more formally tested by looking at the results of first

Figure 8: Loan to value ratios by age and year moved in (1960s cohort)

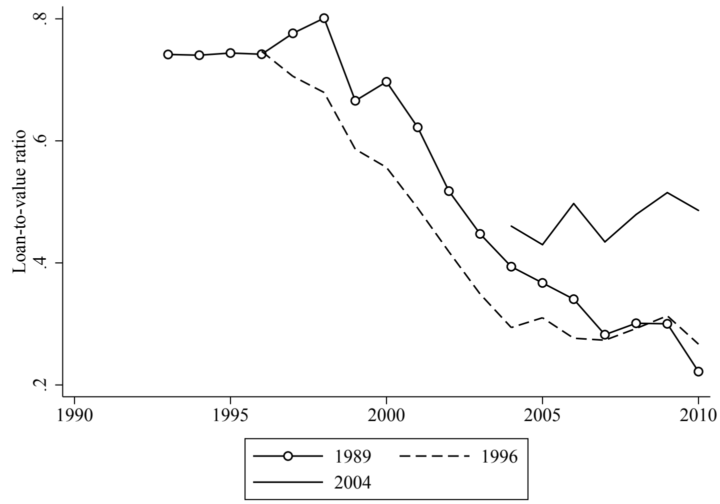

Source: British Household Panel Survey/Understanding Society

stage regressions. We do this in Table 3.

We have two first stage regressions, one for leverage and one for leverage interacted with house prises. In both cases, the F-statistics are greater than the value of 10 suggested as a rule of thumb by Staiger and Stock (1997) for IV estimated using a single sample.[3] Two sample IV methods may suffer less of a bias than standard 2SLS estimators, as errors in the first stage estimation will be unrelated to errors in the second stage equation. This is indeed the rationale for estimators that run first and second stages in split samples (Angrist & Krueger, 1995)). Nonetheless weak instruments may still result in coefficients being biased towards zero in finite samples. The relatively strong first stage we obtain is therefore reassuring. Kleibergen-Paap statistics for the first stage also heavily reject the hypothesis of underidentification.[4]

[3]Later we will also report results for a younger sub-sample of households (those with heads aged 25-45). In this sub-sample, we obtain an even stronger first stage (with F-stats for ω_{t-1} and $\omega_{it-1} \times (\frac{p_{rt}}{p_{rt-1}} - 1)$ of 19.75 and 80.20 respectively).

[4]A further 'first stage' check we can conduct is to look for a positive association between our instrument and total mortgage debt in the LCFS. This would us that the association between our instrument and leverage is not limited to our first sample. Regressing mortgage debt on (P/Y_{-T}) and our controls yields a positive coefficient with a t-statistic of 24.07.

Table 3: First stage results

	ω_{it-1}	$\omega_{it-1} \times (\frac{p_{rt}}{p_{rt-1}} - 1)$
P/Y_{-T}	0.385	-0.010
	(0.042)	(0.003)
$P/Y_{-T} \times (\frac{p_{rt}}{p_{rt-1}} - 1)$	0.510	0.667
	(0.358)	(0.050)
Shea partial R^2	0.007	0.025
F-stat (p-value)	40.01	114.17
	(<0.001)	(<0.001)
Kleibergen-Paap (p-value)	72.82	
	(<0.001)	
N	29,695	
Clusters	7,813	

The second requirement for a suitable instrument is that it is itself uncorrelated with the error term.

In our case, there may be concerns that those who move home in years with higher price-income ratios will have spending patterns that are different to those who moved in other years for reasons other than the degree of their leverage. The most obvious challenge is that since price-income ratios have tended to increase over time, those households with higher values of our instrument will tend to have moved more recently. They may therefore be younger, or be more likely to furnishing a new home. We address these concerns of this nature directly by including a control for years households have spent in their current address (in addition to a dummy variable for households having moved in in the last year to account for first year "setting up" expenses). By including a full set of cohort-region-year interactions we also control for any region or cohort specific trends in income growth that may be correlated with house price changes. Theses dummies can be thought of capturing shocks that are potentially correlated with house price movements but differ in their effects across young and old or across different regional labour markets. One such shock is to future income expectations which would be expected to boost the consumption of younger (and so more leveraged) cohorts by more. If effect such as these are not controlled for they could lead to spuriously large estimates of house price wealth effects for younger households (Attanasio, Blow, Hamilton, & Leicester, 2009).

The use of our instrument in combination with these controls means we effectively compare the spending responses of house price changes between two households in the same region and same cohort but who moved into their homes at different times (when credit was either looser or tighter).

Questions about endogeneity may remain however. For example, households may have been more likely to move when house prices were high because greater unobservable wealth made them less price sensitive. They may also have moved into larger houses. This would create a spurious association between our instrument and consumption. Households who moved at times when credit was loose may also be more likely to move in response to economic shocks and drop out of our sample introducing a selection bias. The assumption that such omitted factors do not induce a correlation between instruments and the error term is usually something which cannot be verified. Omitted variables are typically omitted because they are unobserved. However, when using a two sample approach, such tests are possible. Some variables may be observed in the sample in which we run our first stage regressions even if they are not present in our main sample

To address additional endogeneity concerns we therefore look for an association between our instruments and gross house values, asset incomes and the probability of being a mover in the BHPS and Understanding Society panels conditional on our covariates. Table 4 reports results from regressions of these potential sources of endogeneity on our instruments and our other covariates. The instruments are both jointly and individually insignificant in all models suggesting that they are plausibly orthogonal to these omitted variables. In addition to the results shown in Table 4, we also regress unsecured debt to income ratios and an indicator for whether households have positive debts on our instruments. Debts are only observed in 3 of the 18 waves of the BHPS survey, and so these tests are necessarily conducted on a much smaller sample. The instruments are again individually and jointly insignificant in these regressions.

Table 4: Exogeneity of instruments

Panel a	Credit conditions			
Dependent var.	$\log(HValue)$	Invest inc. > 1000	Invest inc.$= 0$	Mover_{t+1}
P/Y_{-T}	0.012	-0.002	-0.006	-0.001
	(0.012)	(0.009)	(0.014)	(0.004)
$P/Y_{-T} \times (\frac{p_{rt}}{p_{rt-1}} - 1)$	-0.053	-0.064	0.092	0.013
	(0.105)	(0.078)	(0.123)	(0.037)
F-test (p-values)	0.448	0.712	0.658	0.873
N	29,411	28,251	28,251	22,743
Clusters	7,697	7,661	7,661	6,334
Panel b	Lagged leverage			
LTV_{t-1}	-0.202***	-0.167***	0.333***	0.0001
	(0.021)	(0.014)	(0.0220)	(0.008)
$LTV_{t-1} \times (\frac{p_{rt}}{p_{rt-1}} - 1)$	-0.622***	-0.142	0.175	0.049
	(0.199)	(0.138)	(0.218)	(0.0860)
F-test (p-values)	< 0.001	< 0.001	< 0.001	0.847
N	29,411	28,251	28,251	22,743
Clusters	7,697	7,611	7,661	6,334

Note: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses. Controls for education, cohort-region-year dummies, sex, house type, number of rooms, number of adults, number of children, years at address, and a dummy variable for having moved in in the previous year. Standard errors are clustered at the individual level.

A further exercise we can do is test to how our instrument compares to the use of households lagged LTV ratios to sort households according to their leverage. This is the source of variation used in a number of previous studies (e.g. Disney et al. (2010), Dynan (2012)). The results of this comparison are shown in the second panel of Table 4. There is strong evidence that those with higher lagged leverage have fewer financial assets and tend to live in less valuable homes.[5]

[5] As an additional check we regress our instrument P/Y_{-T} and lagged leverage on possible confounding variables simultaneously. We find again that these variables are not correlated with credit conditions when consumers moved into their homes, but are related to households' lagged leverage.

5.2 Results

Table 5 shows results for using our estimation equation (**??**) for different forms of spending. These include total spending, non-durable spending (consumption), durable spending (excluding residential investment), residential investment spending and total spending less residential investment. We log each of these spending variables with the exception of residential investment and durable spending, which we transform using the inverse hyperbolic sine (IHS) transformation as a significant fraction of households record zero spending on these categories. This is defined as

$$IHS(y) = \log(y + \sqrt{(y^2 + 1)}) \tag{17}$$

The IHS transformation approximates log values at high values of spending but remains defined at zero.

The regressions show that leveraged returns are positively associated with both increased total and consumption spending. Our results imply that for a 10% increase in house prices, total spending would be just 0.50 % higher for outright owners and 1.00 % higher for those with a loan-to-value ratio of 50% (i.e ω_{it}=2). The equivalent non-durable consumption responses fall to -0.2% and -0.4% respectively. These responses are small and not statistically significant, suggesting only limited housing wealth effects on consumption in our sample. For durable goods, we observe a negative response to house price increases, though our estimates are imprecise and not significant.

Column (4) shows however that there is strong response for residential investment spending. We estimate that a 10% increase in house prices results in a 6.9% increase in residential investment spending for outright owners, rising to 14% for those with a loan-to-value ratio of 50%.

Table 5: Log spending responses

	Total	Non-durables	Durables (IHS)	Res inv. (IHS)	Total - Res
	(1)	(2)	(3)	(4)	(5)
ω_{it-1}	-0.020	-0.017	-0.042	-0.046	-0.022*
	(0.013)	(0.011)	(0.046)	(0.048)	(0.012)
$\omega_{it-1} \times (\frac{p_{rt}}{p_{rt-1}} - 1)$	0.053	-0.024	-0.189	0.749***	-0.004
	(0.072)	(0.066)	(0.267)	(0.281)	(0.054)
R^2	0.355	0.377	0.113	0.082	0.361
N	60,357	60,357	60,357	60,357	60,357

Note: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses. Controls for education, cohort-region-year dummies, sex, house type, number of rooms, number of adults, number of children, years at address, and a dummy variable for having moved in in the previous year.

In Table 6 we separately consider results for a younger subsample of households (those with heads aged 25-45). The results are similar to those in Table 5 except that if anything residential investment responses are even larger (with an associated elasticity of 0.74). As in the case of our full sample, there is no evidence of a differential response in total spending between leveraged and non-leveraged households.

Table 6: Log spending responses: Ages 25-45

	Total	Non-durables	Durables (IHS)	Res inv. (IHS)	Total - Res
	(1)	(2)	(3)	(4)	(5)
ω_{it-1}	-0.013	-0.003	0.002	-0.105***	-0.001
	(0.010)	(0.009)	(0.037)	(0.041)	(0.009)
$\omega_{it-1} \times (\frac{p_{rt}}{p_{rt-1}} - 1)$	0.056	-0.009	-0.452**	0.831***	-0.003
	(0.049)	(0.044)	(0.190)	(0.207)	(0.047)
R^2	0.319	0.324	0.106	0.092	0.314
N	29,557	29,557	29,557	29,557	29,557

Note: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses. Controls for education, cohort-region-year dummies, sex, house type, number of rooms, number of adults, number of children, years at address, and a dummy variable for having moved in in the previous year.

5.3 Robustness checks

In addition to these results, we carry out a range of robustness checks. For reasons of space, we discuss the results of these briefly here, reporting the full set of results in Appendix B.

1. *Alternative definitions of residential investment.* The definition of residential investment we use above is relatively broad compared to what would for example be used in the national accounts. In particular, we include other fixtures and durable investments (such as for example, kitchen equipment) that we consider likely to be capitalised into the value of the property but which may be excluded in other other definitions. In the appendix, we also consider a narrow definition that is restricted to spending on changes to structure of the property as well as household repairs and maintenance.[6] The results we obtain are very similar to our main results. We also find a positive effect when we use an indicator of whether households made investments in household

[6] This definition is more in the spirit of national accounts. For example in the US National Income and Product Accounts, Private Fixed Investment includes “construction of new nonresidential and residential buildings.”, “improvements (additions, alterations, and major structural replacements) to nonresidential and residential buildings.” and “certain types of equipment (such as plumbing and heating systems and elevators) that are considered an integral part of the structure.” (see Chapter 6 of U.S. Bureau of Economic Analysis (BEA) (2016))

extensions specifically as our dependent variable.

2. *Alternative instruments.* We also consider results using two alternative instruments. The first is the Credit Conditions Index used in Fernandez-Corugedo and Muellbauer (2006), and the second is the average house price in each region at the time individuals moved into their current homes (as used as an instrument for mortgage debt in Chetty, Sandor, and Szeidl (2017)). The use of these alternative instruments do not give very different results.

3. *Sample restrictions.* We exclude households who moved within the previous year from our analysis, but concerns may remain that our spending effects are driven by more recent movers (who are likely to be the most levered, and possibly at credit a constraint). We therefore consider results from an alternative sample where we exclude those who moved into their homes within the previous five years. Within this subsample we find slightly larger effects of house price changes on the investment spending of households who are predicted to be the most levered.

6 Conclusion

In this paper, we show that viewing leverage as a portfolio choice over risky assets can help to explain non-linear effects of leverage in amplifying wealth shocks on consumption, deleveraging and re-leveraging behaviour over the business cycle and also residential investment over the house price cycle.

Empirically we find evidence that households view housing as a financial asset. In response to house price increases, leveraged households borrow and invest more in housing in order to rebalance their portfolios.

Our results suggest that leverage works to amplify the effects of house price shocks. Moreover house price rises themselves lead households to re-leverage and so to increase the 'size' of their balance sheets. These effects would be expected to in turn make households more sensitive to future shocks to their housing wealth.

These findings suggest that macro-prudential policy interventions aimed at limiting re-leveraging behaviour could have an important role in stablising spending over the business cycle.

Appendix A Two-sample IV

In this paper we make use of Two Sample Two Stage Least Squares (TS2SLS). Inoue and Solon (2010) show that this approach is more efficient than the TSIV estimator of Angrist and Krueger (1992).

TS2SLS is best explained by first considering a standard two-stage least squares (2SLS) approach.

Let $\mathbf{M} = [\ X \quad \omega_{t-1} \quad \omega_{t-1} \times (\frac{p_{rt}}{p_{rt-1}} - 1)\ ]$ denote the $n \times (k+p)$ matrix of right-hand side variables (p of which are endogenous). Suppose we face the problem of consistently estimating the $1 \times (k+p)$ vector of coefficients δ in the model

$$c = \mathbf{M}\delta + e$$

where ω_{t-1} and e are correlated. It is well known that the coefficients estimated using a naive OLS regression of c on $\mathbf{M}$ will be biased. To solve this problem, instrumental variable methods make use of an $n \times (k+q)$ matrix of instruments $\mathbf{Z}$ where the p endogenous variables in $\mathbf{M}$ are replaced with $q \geq p$ variables that are assumed to be exogenous. This assumption implies that $E[e|\mathbf{Z}] = 0$ and means that δ can be consistently estimated using the 2SLS estimator

$$\hat{\delta}_{2SLS} = (\hat{\mathbf{M}}'\hat{\mathbf{M}})^{-1}\hat{\mathbf{M}}'c \tag{18}$$

where $\hat{\mathbf{M}} = \mathbf{Z}(\mathbf{Z}'\mathbf{Z})^{-1}\mathbf{Z}'\mathbf{M}$, or the fitted values from the set of reduced form regressions of the columns of $\mathbf{M}$ on $\mathbf{Z}$

$$\mathbf{M} = \mathbf{Z}\Pi + v$$

Notice here that while this estimator requires knowledge of both the cross-products $\mathbf{Z}'\mathbf{M}$ and $\mathbf{Z}'c$ we do not require the cross product $\mathbf{M}'c$. This insight was the basis for two sample IV proposed in Angrist and Krueger (1992).[7] They show that under certain conditions, it is possible to estimate δ even if no sample can be found that contains data on $\mathbf{M}$, c and $\mathbf{Z}$ simultaneously. All that is required is a sample that includes both c and $\mathbf{Z}$ (but not necessarily the endogenous components of $\mathbf{M}$) and another which includes $\mathbf{Z}$ and $\mathbf{M}$ (but not necessarily c) . This allows us to calculate a two sample 2SLS estimator (TS2SLS) that is analagous to (18)

$$\hat{\delta}_{TS2SLS} = (\hat{\mathbf{M}}_1{}'\hat{\mathbf{M}}_1)^{-1}\hat{\mathbf{M}}_1{}'c_1 \tag{19}$$

where $\hat{\mathbf{M}}_\mathbf{1} = \mathbf{Z_1}(\mathbf{Z_2'Z_2})^{-1}\mathbf{Z_2'M_2} = \mathbf{Z}_1\hat{\Pi}_2$. Here c_1 and $\mathbf{M}_1$ contain n_1 observations from the first sample while $\mathbf{M}_2$ and $\mathbf{Z}_2$ contain n_2 observations from the second. $\hat{\Pi}_2$ is the coefficient matrix formed from a regression of $\mathbf{M}_2$ on $\mathbf{Z}_2$.

This estimator can be implemented using a simple two step procedure:

[7] In their orginal article, Angrist and Krueger (1992) in fact proposed orginally an alternative GMM estimator $\hat{\delta}_{IV} = (\mathbf{Z}_2'\mathbf{M}_2/n_2)^{-1}(\mathbf{Z}_1'c_1/n_1)$. Asymptotically this gives identical results to the TS2SLS estimator. However, Inoue and Solon (2010) show these two approaches will in general give different answers in finite samples, and that the TS2SLS is more efficient. This gain in efficiency arises because the latter estimator corrects for differences in the two samples in the distribution of $\mathbf{Z}$

1. Run a first stage regression in sample 2 and using the recovered coefficients to impute $\mathbf{M}$ in sample 1.

2. In sample 1, regress c_1 on the imputed values of $\mathbf{M}$ to recover $\hat{\hat{\delta}}_{TS2SLS}$.

We adjust standard errors from our second stage regression to account for the two-step nature of the procedure. Because we cluster observations from the same household in our first stage regression, we use the robust standard error correction for TS2SLS derived in Pacini and Windmeijer (2016).

Appendix B Alternative estimation approaches

B.1 Alternative definitions of residential investment

First we investigate the extent to which our results depend on our chosen measure of residential investment. The measure of residential investment that we use for our main results includes certain white goods such as cookers, refrigerators and washing machines which are often capitalised into property values but which would not necessarily be considered residential investment spending in for instance a national accounting framework. Here we examine the extent to which our results are robust to the removal of these items by restricting out definition to goods such as electric tools, floor coverings and the costs of installing or repairing heating and air conditioning units (along with spending on household extensions).

We show results using these alternative measures in Table 7. Column (1) shows results using the inverse hyperbolic sine transformation of our narrower residential investment measure. The effects of increases in prices for more leveraged households are still large and statistically significant (and indeed very similar to those obtained in our main results). In Column (2) we show results from a linear probability model in which the dependent variable takes a value of 1 if the household is observed spending a positive amount on household extensions. This is probably the purest measure of residential investment in that it only includes structural modifications to the home. Again we find that the investment spending of more leveraged households is significantly more responsive to house price changes than the spending of other home-owners. Each 10% increase in house prices associated with a 1% increase in the probability that a household builds an extension.

Table 7: Results with alternative definitions of residential investment

	Narrow Res inv.	Extensions>0
	(1)	(2)
ω_{it-1}	-0.046	-0.011
	(0.048)	(0.010)
$\omega_{it-1} \times (\frac{p_{rt}}{p_{rt-1}} - 1)$	0.734***	0.102*
	(0.280)	(0.057)
R^2	0.081	0.039
N	60,357	60,357

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses. See table 5 for list of controls.

B.2 Alternative instruments

In this section, we consider how our results are affected when we use two alternative instruments in place of the price-to-income ratio at the time individuals moved into their current residences.

The first of these is the Credit Conditions Index (CCI) assembled in Fernandez-Corugedo and Muellbauer (2006). This index contains 10 indicators of credit conditions. Two are aggregate measures of unsecured and mortgage debts. The remaining 8 are fractions of mortgages for first time buyers that are above given loan-to-value and loan-to-income ratios for different age groups and regions. The index is constructed controlling for various determinants of credit demand to ensure the index reflects credit supply conditions.[8] The series is plotted alongside our instrument in Figure 9. The CCI shows a discontinuous increase in 1981. Because this is not matched by a similarly discontinuous increase in leverage for those moving in these years in our sample, when we include move years before this date we find the instrument to be weak and our results imprecise. The first two columns of Table 8 present results for log total spending and residential investment (conditional on moving in 1981 or after). The results are very similar to what we obtain in our main specification, with the implied elasticity much greater for residential than other forms of spending.

[8]These controls are: nominal and real interest rates, a measure of interest rate expectations and of inflation and interest rate volatility, mortgage and housing return, 36 risk indicators, house prices, income, a proxy for expected income growth, the change in the unemployment rate, demography, consumer confidence, portfolio wealth components, proxies for sample selection bias and institutional features.

Figure 9: Credit Conditions Index vs price-income ratio

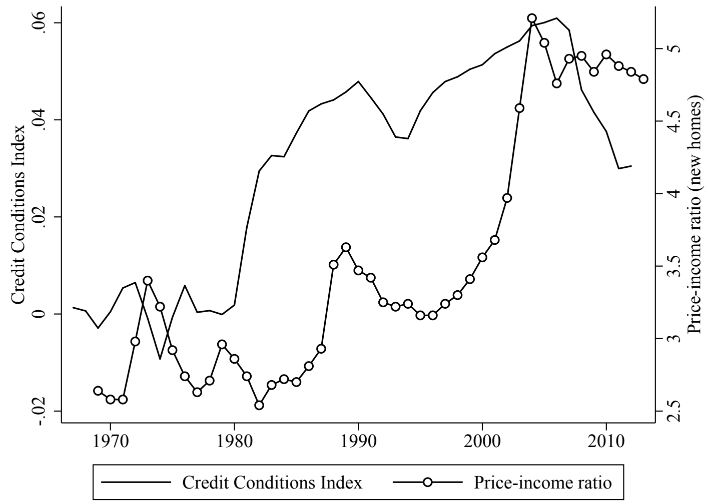

The second alternative instrument we consider is the average regional price at the point homeowners moved into their homes. This makes use of interregional variation as well as intertemporal variation in house prices. We report results for this approach in Table 8. We find that they are again very similar to our main results.

Table 8: Results with alternative instruments

	CCI		Reg. house prices	
	Total	Res. inv	Total	Res. inv
	(1)	(2)	(3)	(4)
ω_{it-1}	-0.030	-0.152	0.034**	0.104*
	(0.041)	(0.165)	(0.015)	(0.059)
$\omega_{it-1} \times (\frac{p_{rt}}{p_{rt-1}} - 1)$	0.095	0.847*	0.153*	0.824***
	(0.115)	(0.451)	(0.072)	(0.317)
Instruments:				
CCI_{-T} and $CCI_{-T} \times (\frac{p_{rt}}{p_{rt-1}} - 1)$	x	x		
P_{-rT} and $P_{-rT} \times (\frac{p_{rt}}{p_{rt-1}} - 1)$			x	x
R^2	0.353	0.087	0.355	0.082
N	52,155	52,155	60,357	60,357

Notes: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses.

See table 5 for list of controls.

B.3 Alternative sample

A further concern might that our results for more leveraged households are driven entirely by households who have just moved into their homes (and are thus more likely to be at a credit constraint). Since price-to-income ratios have tended to increase over time, our first stage regressions will tend to predict higher rates of leverage for more recent movers.

In our main results, we exclude households who moved into their homes in the previous year only. In Table 9 we consider how our results are affected when we exclude households who moved into their homes within the previous five years. The results from this exercise are remarkably similar to our main set of results.

Table 9: Log spending responses

	Total	Non-durables	Durables (IHS)	Res inv. (IHS)	Total - Res
	(1)	(2)	(3)	(4)	(5)
ω_{it-1}	-0.026**	-0.016	-0.054	-0.065	-0.024*
	(0.013)	(0.012)	(0.049)	(0.051)	(0.013)
$\omega_{it-1} \times (\frac{p_{rt}}{p_{rt-1}} - 1)$	-0.004	0.000	-0.315	0.907**	-0.012
	(0.120)	(0.107)	(0.444)	(0.458)	(0.116)
R^2	0.378	0.405	0.122	0.079	0.387
N	42,285	42,285	42,285	42,285	42,285

Note: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses. See table 5 for list of controls.

References

Adrian, T., & Shin, H. S. (2010). Liquidiy and leverage. *Journal of Financial Intermediation*, *19*, 418-437.

Aladangady, A. (2017). Housing wealth and consumption: Evidence from geographically linked microdata. *American Economic Review*, *107*(11), 3415-3446.

Angrist, J. D., & Krueger, A. B. (1992). The effect of age at school entry on educational attainment: An application of instrumental variables with moments from two samples. *Journal of the American Statistical Association*, *87*, 328-336.

Angrist, J. D., & Krueger, A. B. (1995). Split sample instrumental variables. *Journal of Business and Economic Statistics*, *13*(2), 225-235.

Attanasio, O. P., Blow, L., Hamilton, R., & Leicester, A. (2009). Booms and busts: Consumption, house prices and expectations. *Economica*, *76*, 20-50.

Bénétrix, A., Eichengreen, B., & O'Rourke, K. (2012). How housing slumps end. *Economic Policy*, *72*, 647-692.

Berger, D., Guerrieri, V., Lorenzoni, G., & Vavra, J. (2017). House prices and consumer spending. *Review of Economic Studies*, *Forthcoming*.

Browning, M., Crossley, T. F., & Winter, J. (2014). The measurement of household consumption expenditures. *Annual Review of Economics*, *6*, 475-501.

Chetty, R., Sandor, L., & Szeidl, A. (2017). The effect of housing portfolio choice. *Journal of Finance*, *72*(3), 1171-1212.

Christelis, D., Georgarakos, D., & Japelli, T. (2015). Wealth shocks, unemployment shocks and consumption in the wake of the great recession. *Journal of Monetary Economics*, *72*, 21-41.

Cloyne, J., Huber, K., Ilzetzki, E., & Kleven, H. (2017). The effect of house prices on household borrowing: A new approach. *Working Paper*.

Cooper, D. (2010). Did easy credit lead to overspending? home equity borrowing and household behavior in the early 2000s. *Federal Reserve Bank of Boston Public Policy Discussion Papers 09-7*.

Cooper, D. (2013). House price fluctuations: the role of housing wealth as borrowing collateral. *The Review of Economics and Statistics*, *95*(4), 1183-1197.

Disney, R., Bridges, S., & Gathergood, J. (2010). House price shocks and household indebtedness in the united kingdom. *Economica*, *77*, 472-496.

Disney, R., & Gathergood, J. (2011). House price growth, collateral constraints and the accumulation of homeowner debt in the united states. *The B.E. Journal of Macroeconomics*, *11*(1), 1-30.

Disney, R., Gathergood, J., & Henley, A. (2010). House price shocks, negative equity, and household consumption in the united kingdom. *Journal of the European Economics Association*, *8*(6), 1179-1207.

Dynan, K. (2012). Is a household debt overhang holding back consumption. *Brookings Papers on Economic Activity*.

Fernandez-Corugedo, E., & Muellbauer, J. (2006). Consumer credit conditions in the united kingdom. *Bank of England working papers 314*.

Ghent, A. C., & Kudlyak, M. (2011, September). Recourse and residential mortgage default: Evidence from us states. *Review of Financial Studies*, *24*(9), 31393186.

Henderschott, P. H., Pryce, G., & White, M. (2003). Household leverage and the deductibility of home mortgage interest: Evidence from u.k house purchasers. *Journal of Housing Research*, *14*(1), 49-82.

Inoue, A., & Solon, G. (2010). Two-sample instrumental variables estimators. *The Review of Economics and Statistics*, *92*(3), 557-561.

Lehnert, A. (2004). Housing, consumption, and credit constraints. *Board of Governers of the Federal Reserve System of Finance and Economis Discussion Paper*.

Merton, R. (1969). Lifetime portfolio selection under uncertainty: The continuous-time case. *The Review of Economics and Statistics,*, *51*(3), 247-257.

Mian, A., Rao, K., & Sufi, A. (2013). Household balance sheets, consumption, and the economic slump. *Quarterly Journal of Economics*, *128*(4), 1687-1726.

Mian, A., & Sufi, A. (2011). House price, home-equity-based borrowing, and the us household leverage crisis. *American Economic Review*, *101*(5), 2132-2156.

Pacini, D., & Windmeijer, F. (2016). Robust inference for the two-sample 2sls estimator. *Economics Letters*, *146*, 50-54.

Staiger, D., & Stock, J. (1997). Instrumental variables regression with weak instruments. *Econometrica*, *65*(3), 557-586.

University of Essex. Institute for Social and Economic Research. (2010). British household panel survey: Waves 1-18, 1991-2009. *7th Edition. UK Data Service. SN: 5151*. Retrieved from `http://doi.org/10.5255/UKDA-SN-5151-1`

University of Essex. Institute for Social and Economic Research. (2016). Understanding society: Waves 1-6, 2009-2015. *8th Edition. UK Data Service. SN: 6614*.

U.S. Bureau of Economic Analysis (BEA). (2016). *Concepts and methods of the U.S. national income and product accounts*. Washington, D.C: BEA.