How large are leverage effects?: investigating the link between consumption and leverage using two samples

Thomas F. Crossley and Peter Levell*

February 15, 2016

Abstract

In this paper we estimate the degree to which leverage amplifies the effects of house price shocks on consumer spending. We do so by using instrumental variable methods that are able to combine the information in two datasets. The first is a panel with information on household balance sheets which does not include consumption spending. The other is a survey with detailed consumption data that does not include information on wealth. As well as applying standard two sample IV methods, we show how these can be extended to estimate dynamic relationships between leverage and consumption changes despite only observing repeated cross-sectional data on consumption. We find evidence that each 10% increase in leverage increases the size of housing wealth effects by roughly 10%. Effects are larger for durable spending.

Keywords: House prices, leverage, consumption

[PRELIMNARY AND INCOMPLETE- DO NOT CITE]

1 Introduction

The recent economic experience of the UK has much in common with that of many other developed countries. In the years running up to the 2008 Great Recession, debt to income ratios reached historic highs of over 160% of GDP (Bunn and Rostom, 2014) as consumption spending boomed. This was then followed by a prolonged slump in consumer spending. This pre-crisis increase in debt was partly the result of an increased tendency to extract equity from homes. Figure 1 shows housing equity withdrawal - a measure of the change

*Crossley: Department of Economics, University of Essex Wivenhoe Park, Colchester, UK, C04 3SQ (email: tcross@essex.ac.uk). Levell: Institute for Fiscal Studies, 7 Ridgmount Street, London WC1E 7AE (email: peter_l@ifs.org.uk). The authors would like to thank Laura Blow, Cormac O'Dea and William Elming 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).

in housing equity caused by changes in the stock of debt (from taking out new debt or repaying existing debt) improvements made to existing properties, but excluding changes due to revaluations of home values - over the period 1992-2014. This reached levels of around 5% in 2007 before becoming negative in the years following the financial crisis.[1] At the same time, household loan to value ratios fluctuated wildely as figure 2 shows. Mortgages did not increase as fast as house prices in the years prior to the financial crisis but, after the 2008 crisis household leverage increased sharply as prices fell.

This paper considers what role leverage may have played in the accentuating pre-crisis increases in consumption and in the subsequent consumption bust. This is a question of high policy relevance. Several countries are considering adopting macro-prudential regulations which would seek to limit the growth of leverage precisely in order to dampen economic cycles. Judicious use of such policies requires a thorough understanding of the magnitudes and mechanisms by which leverage may amplify consumption responses to economic shocks.

Figure 1: Housing equity withdrawal

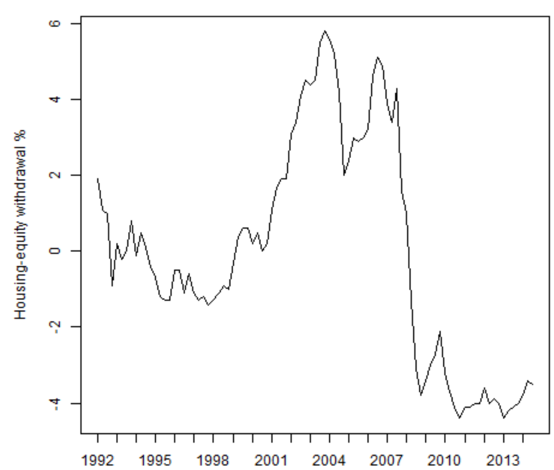

Source: Bank of England

[1] See Reinhold (2011) for explanations of this later development.

Figure 2: Leverage

Source: BHPS and Understanding Society

To answer these questions requires microeconomic data on both spending and household balance sheets. Unfortunately, this information is rarely contained in the same datasets. This has meant previous studies into this question have tended to make use of proxies for consumption (Lehnert, 2004), measures of certain types of savings (Disney, Henley and Gathergood, 2010), or changes in household indebtedness (Mian and Sufi, 2011, Disney and Gathergood, 2011). It is not often obvious how changes in these variables map into changes in spending. In addition, we argue that an understanding of the nature and composition of consumption responses can help us to seperate out the particular channels through which leverage may operate.

The contributions of this paper are twofold. The first contribution is to apply two sample methods of the kind proposed in Angrist and Krueger (1992) to estimate the degree to which leverage amplifies the effects of house price shocks. We use an instrumental variable present in both datasets to combine information on household wealth contained in a long-running panel survey with data in a detailed cross-sectional survey of household spending. This allows us to treat leverage as an endogenous variable that we predict in a first stage run in one sample, and then effectively impute as a regressor in our second sample. We are then able to estimate how the responses of more levered households to house price changes differ across different types of spending. The second, methodological contribution is to extend two sample methods to estimate dynamic relationships between leverage and spending changes despite only having repeated cross-sectional data on consumption.

The rest of this paper is structured as follows. In section 2 we give thoughts on why we think leverage should affect consumer responses to house price changes. Section 3 outlines established two sample methods and a proposed "three sample" extension. In section 4 we discuss our data and estimation. Section 5 discusses our results and relates them to the prior literature.

2 Theoretical motivation

This paper seeks to empirically measure the extent to which leverage affects consumers' responses to changes in house prices. In order to interpret our findings it is worth discussing what mechanisms might generate different responses as house prices rise and fall. This is perhaps best illustrated by means of a thought experiment involving two homeowners: A and B. A and B are identical in terms of their characteristics (age, income, family size) and preferences, and have homes and financial assets of equal value. However, suppose A has a larger mortgage than B for some exogenous reason. If house prices rise, there are a variety of reasons why A's spending may respond differently to B's. The first is that proportionally, A has experienced a larger wealth increase. If both consumers owned homes with a value of £100,000, and A had a 90% mortgage and B a 20% mortgage, then a 10% increase in house prices would represent a doubling of A's housing wealth compared to a 25% increase in B's. This mechanism is usually referred to as a *wealth channel* for the effects leverage. A second reason is that A's leverage position may indicate that she was credit constrained. In this case, prior to the house price increase A was not consuming as much as she would like to given her expected path of future incomes, while B was smoothing her consumption optimally. A would then tend to consume more out of an increase in housing wealth than B. To do so she would have to extract equity from her home by for example increasing her existing mortgage. More subtly, by offering her now more valuable home as collateral she may also be able to borrow more against her future earnings (allowing her to borrow more than the increase in her home value). These mechanisms are sometimes referred to as the credit or *collateral channel*.

The relative strengths of these different channels could depend on a number of factors. The strength of the collateral channel will depend for instance on the degree to which increased housing wealth can be accessed when it is needed. Fixed costs of extracting equity would tend to reduce the strength of this channel. The determinants of the strength of the wealth channel are likely to much more complex. For example it can be argued that wealth effects in general should be small or zero across all consumers. This is because a shock to house prices increases the resources currently available to homeowners but also increases the price of future housing services. These two effects exactly offset each other in some models with infinitely-lived agents in which house price changes leave the present value of expected net wealth unchanged (e.g. Sinai and Souleles, 2005). On the other hand, Berger et al. (2015) present a model where housing wealth effects arise as a result of borrowing constraints and uncertainty. They calibrate their model and find that these effects are large.

To clarify exactly which features of the consumers may lead to a larger or smaller wealth channel among leveraged households, we set up and solve the following simple model.

Households maximise the expected utility function

$$U = E_t\left[\sum_{t=0}^{T}\beta^t u(C_t, H_t)\right]$$

where C_t denotes non-housing consumption, H_t represents the stock of housing owned by the consumer and β is a constant discount factor. Within-period preferences are Cobb-Douglas

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

with the parameter ρ determining the consumer's desire to substitute intertemporally and their risk aversion.

Housing is bought and sold at a cost p_t and depreciates according to the rate δ from one period to the next. Because it is durable, purchases of housing is a form of investment as well as consumption. In addition to housing, households may also purchase a risk free asset a_t which earns an interest rate r. Total wealth in period t is thus $W_t = (1+r)a_{t-1} + p_t(1-\delta)h_{t-1}$.

The per period budget constraint is

$$C_t + p_t H_t = Y_t + (1+r_t)(W_{t-1} - C_{t-1} - p_{t-1}H_{t-1}) + p_t(1-\delta)H_{t-1}$$

Income follows the process

$$Y_t = G_t P_t T_t$$

where G_t is a deterministic trend, T_t is a transitory shock and P_t a permanent shock which evolves according to

$$P_t = P_{t-1}\Psi_t$$

Both T_t and Ψ_t are lognormally distributed and have mean 1. House prices evolve according to a random walk. In addition consumers face a borrowing constraint $a_t > 0$.

This model cannot be solved analytically. Instead we use dynamic programming to solve the following recursive problem

$$V_t(W_t, s_t) = \max_{C_t, H_t} U(C_t, H_t) + \beta E_t[V_{t+1}(W_{t+1}, s_{t+1})]$$

where s_t denotes the state of the world (current level of permanent income and house price). We then plot the policy function for consumption against liquid assets a_t for individuals with the same level of housing stock but in two different price regimes.

[GRAPH SHOWING DIFFERENT POLICY FUNCTIONS IN DIFFERENT REGIMES OF RISK]

3 Instrumental variable methods

3.1 Two sample IV

Suppose we face the problem of consistently estimating the $1 \times (k+p)$ coefficient vector β in a model of the form

$$Y = \mathbf{X}\beta + \varepsilon$$

where ε is a mean zero error and $\mathbf{X}$ is an $n \times (k+p)$ matrix of which p variables are correlated with ε. It is well known that the coefficients estimated using a naive OLS regression of Y on $\mathbf{X}$ 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 X are replaced with $q \geq p$ variables that are assumed to be exogenous. This assumption implies that $E[\varepsilon|\mathbf{Z}] = 0$ and means thatβ can be consistently estimated using the two-stage least squares (2SLS) estimator

$$\hat{\beta}_{2SLS} = (\hat{\mathbf{X}}'\hat{\mathbf{X}})^{-1}\hat{\mathbf{X}}'Y \tag{1}$$

where $\hat{\mathbf{X}} = \mathbf{Z}(\mathbf{Z}'\mathbf{Z})^{-1}\mathbf{Z}'\mathbf{X}$. These are the fitted values from reduced form regressions of $\mathbf{X}$ on $\mathbf{Z}$

$$\mathbf{X} = \mathbf{Z}\Pi + \eta$$

Notice here that while this estimator requires knowledge of both $\mathbf{Z}'\mathbf{X}$ and $\mathbf{Z}'Y$ we do not require the cross product $\mathbf{X}'\mathbf{Y}$. This insight was the basis for two sample IV proposed in Angrist and Krueger (1992). They show that under certain conditions, it is possible to estimate β even if no sample can be found that contains data on $\mathbf{X}$, Y and $\mathbf{Z}$ simultaneously. All that is required is a sample that includes both Y and $\mathbf{Z}$ (but not necessarily $\mathbf{X}$) and another which includes $\mathbf{Z}$ and $\mathbf{X}$ but not Y. This allows us to calculate a two sample 2SLS estimator (TS2SLS) that is analagous to (1)

$$\hat{\beta}_{TS2SLS} = (\hat{\mathbf{X}}_1{}'\hat{\mathbf{X}}_1)^{-1}\hat{\mathbf{X}}_1{}'Y_1 \tag{2}$$

where $\hat{\mathbf{X}}_{\mathbf{1}} = \mathbf{Z}_{\mathbf{1}}(\mathbf{Z}_{\mathbf{2}}'\mathbf{Z}_{\mathbf{2}})^{-1}\mathbf{Z}_{\mathbf{2}}'\mathbf{X}_{\mathbf{2}} = \mathbf{Z}_1\hat{\Pi}_2$. Here Y_1 and $\mathbf{X}_1$ contain n_1 observations from the first sample while $\mathbf{X}_2$ and $\mathbf{Z}_2$ contain n_2 observations from the second. $\hat{\Pi}_2$ is the coefficient matrix formed from a regression of $\mathbf{X}_2$ on $\mathbf{Z}_2$. This estimator can be implemented by running a first stage regression in sample 2 and using the recovered coefficients to effectively impute $\mathbf{X}$ in the second sample. Standard errors should then be adjusted to correct for the two stage nature of the procedure (Murphy and Topel, 1985). Inoue and Solon (2010) calculate the correction needed for the two sample estimator as

$$1 + [\hat{\beta}'_{TS2SLS}\hat{\mathbf{\Sigma}}_\eta\hat{\beta}_{TS2SLS}/\hat{\sigma}_{11}]$$

where $\hat{\mathbf{\Sigma}}_\eta$ is a consistent estimate of covariance matrix for the first stage errors and $\hat{\sigma}_{11}$ is the variance of residuals from the second stage.[2]

A number of papers have demonstrated that 2SLS is biased in finite samples (Nagar (1959), Phillips and Hale (1977), Bound, Jaeger, and Baker (1995), Staiger and Stock (1997)) towards OLS estimates. This bias arises because the *estimated* coefficients are used to construct fitted values. Letting $\mathbf{P}_{z,i} = \mathbf{Z}_i(\mathbf{Z}_i'\mathbf{Z}_i)^{-1}\mathbf{Z}_i'$ where $\mathbf{Z}_i$ is the ith row of $\mathbf{Z}$, the fitted values for individual i are

$$\mathbf{P}_{z,i}\mathbf{X}_i = \mathbf{Z}_i\Pi + \mathbf{P}_{z,i}\eta_i$$

$\mathbf{P}_{z,i}\eta_i$ will in general be correlated with ε_i unless $\mathbf{X}_i$ and ε_i are uncorrelated (eliminating the need to instrument). The bias is greater when the R^2 of the first stage regression is low and if many instruments are used relative to the number of endogeous variables. Various methods have been proposed to address this bias by breaking the link between η_i and ε_i. These include Jackknife Instrumental Variables Estimator (JIVE) (Phillips and Hale (1977), Angrist, Imbens and Krueger (1999)) and the split-sample estimator (Angrist and Krueger (1995)) which construct fitted values for each individual with coefficients estimated using a sample of individuals excluding i. In the split sample case this is achieved by randomly dividing a single sample in which all three of $\mathbf{X}$, Y and $\mathbf{Z}$ are observed into two, and running the first stage regressions in one subsample and the second stage regression in the other. This solution is inherent to the two-sample estimator as the first and second stages of the process are by definition run on separate samples. However, Angrist and Krueger (1995) show that their split-sample estimator, and so analagously the two sample estimator, are subject to a separate source of bias to 2SLS which attenuates estimates towards zero in finite samples. This bias is given by

$$\theta = \left[X_2 Z_2(Z_2'Z_2)^{-1}Z_1'Z_1(Z_2'Z_2)^{-1}Z_2'X_2\right]^{-1} \times \left[X_2'Z_2(Z_2'Z_2)^{-1}Z_1'X_1\right]$$

which can be approximated in the split sample case by a regression of the fitted values $\hat{\mathbf{X}}_1$ estimated from a second sample and the values of $\mathbf{X}_1$ in the first sample. Angrist and Krueger (1995) use an estimate of θ to correct split sample estimates for bias. Unfortunately this cannot be applied unless $\mathbf{X}$ is observed in both samples. As in the case of 2SLS, this bias is worsened if many weak instruments are used in the first stage. This means that concerns about the first stage being "over-fitted" are just as relevant in the two sample case as in standard IV applications.

3.2 A three sample IV estimator

It is also common for researchers to wish to estimate a model of the form

[2] In their orginal article, Angrist and Krueger (1992) in fact proposed orginally an alternative GMM estimator $\hat{\beta}_{IV} = \left(\mathbf{Z}_2'\mathbf{X}_2/n_2\right)^{-1}(\mathbf{Z}_1'Y_1/n_1)$. 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 $\mathbf{Z}$.

$$Y = \mathbf{X}\beta + \alpha + \varepsilon$$

where α is a vector of unobserved, individual fixed effects that is not only correlated with $\mathbf{X}$ but potentially with $\mathbf{Z}$ as well. In particular let

$$\alpha = f(\mathbf{Z}) + v$$

where $E[v|\mathbf{Z}] = 0$. As before let ε be correlated with an $n \times p$ submatrix of variables in $\mathbf{X}$ (denoted $\tilde{\mathbf{X}}$), but not with $\mathbf{Z}$. With a panel dataset contiaining all variables, α could be eliminated using first differences. Using ΔW_t to denote $W_t - W_{t-1}$ this would imply a model

$$\Delta Y_t = \mathbf{\Delta X}_t \beta + \Delta\varepsilon_t \tag{3}$$

where $\Delta\mathbf{X}_t$ could be instrumented using $\mathbf{Z}$ (or $\Delta\mathbf{Z}$ as appropriate). Suppose however that, as above, $\mathbf{Z}$ and $\mathbf{Y}$ are the only variables observed in one sample (sample 1) while $\mathbf{X}$ and $\mathbf{Z}$ are only variables observed in a second sample (sample 2). If sample 1 were cross-sectional, then a strategy of taking first differences would not open to us since we only observe Y for each individual once.

One way forward in this situation is to note that the absence of lagged information on a particular individual is not to dissimilar to the problem we faced above. With a third sample with data on Y_{t-1} and $\mathbf{Z}$, we could impute it to the first sample in the same way as the endogenous variables in the two sample approach above. In particular let there be a third sample containing data on Y_{t-1} with

$$Y_{3,t-1} = \mathbf{Z}_{3,t-1}\Omega + \eta_{3,t-1}$$

where $E[\eta_{3,t}|\mathbf{Z}] = 0$. This allows us to impute a value of ΔY_t

$$\widehat{\Delta Y_t} = Y_{1,t} - \mathbf{Z}_{1,t-1}(\mathbf{Z}'_{3,t-1}\mathbf{Z}_{3,t-1})^{-1}\mathbf{Z}'_{3,t-1}Y_{3,t-1}$$

$$= Y_{1,t} - \mathbf{Z}_{1,t-1}\Omega - \mathbf{Z}_{1,t-1}(\mathbf{Z}'_{3,t-1}\mathbf{Z}_{3,t-1})^{-1}\mathbf{Z}'_{3,t-1}\eta_{3,t-1}$$

Here $\mathbf{Z}_{1,t-1}$ is of course unobserved for individuals in sample 1 at time t. However, we can set it equal to $\mathbf{Z}_{1,t}$ for most control variables. For example, the value of Y_{t-1}imputed to an individual with two children in period t could be made conditional on having two children in period t-1. This controls for the effect of having children by holding them constant over the period of comparison. Other variables that vary deterministically over time (such as year dummies or age) are known for time $t-1$ given their values at time t and their values in the matrix $\mathbf{Z}_{1,t-1}$ can be set accordingly.

Now note that if we did observe lagged values in sample 1, they would take the form

$$Y_{1,t-1} = \mathbf{Z}_{1,t-1}\Omega_{t-1} + \eta_{1,t-1}$$

Implying that

$$\widehat{\Delta Y_t} = Y_{1,t} - Y_{1,t-1} + u_t$$

where $u = -(\eta_{1,t-1} + \mathbf{Z}_{1,t-1}(\mathbf{Z}'_{3,t-1}\mathbf{Z}_{3,t-1})^{-1}\mathbf{Z}'_{3,t-1}\eta_{3,t-1})$. Using (3) gives

$$\widehat{\Delta Y_t} = \Delta\tilde{\mathbf{X}}_t\beta + \varepsilon + u$$

where only the p endogenous components of $\mathbf{X}$ remain on the right hand side here.[3] This is because the exogenous elements of $\mathbf{X}$ (which are also included in $\mathbf{Z}$) are held constant between t and $t-1$, and are thus differenced out.

When $\mathbf{Z}_1, \mathbf{Y}_1$ and $\mathbf{Z}_2, \mathbf{Y}_2$ are jointly independent, $\mathbf{Z}$ will now be uncorrelated with the error term as

$$E[u_1|\mathbf{Z}_1] == E[\epsilon_1|\mathbf{Z}_1] - E[\eta_1|\mathbf{Z}_1] - \mathbf{Z}_1 \times E[(\mathbf{Z}'_3\mathbf{Z}_3)^{-1}\mathbf{Z}'_3\eta_3]$$

which iterating expectations gives

$$= -\mathbf{Z}_1 \times E\left\{E[(\mathbf{Z}'_3\mathbf{Z}_3)^{-1}\mathbf{Z}'_3\eta_3]|\mathbf{Z}_3\right\} = 0$$

This allows us to estimate β using a three sample estimator

$$\hat{\beta}_{3S2SLS} = (\widehat{\Delta\tilde{\mathbf{X}}}'\widehat{\Delta\tilde{\mathbf{X}}})^{-1}\widehat{\Delta\tilde{\mathbf{X}}}'\hat{Y}$$

where $\widehat{\Delta\mathbf{X}} = \mathbf{Z_1}(\mathbf{Z'_2Z_2})^{-1}\mathbf{Z'_2}(\mathbf{X}_{2,t} - \mathbf{X}_{2,t-1})$. In the Appendix we show that this estimator is subject to an attenuation bias in the same way as the two sample estimator, but that it is also consistent. However, the attenuation bias depends only on how well we are able to fit the values of $\widehat{\Delta\mathbf{X}}$ and not of Y. As before this can be estimated by running OLS regressions in each of our three samples. If we follow this procedure, a correction would need to be applied to standard errors estimated at the final stage. This is not to dissimilar to the correction given by Inoue and Solon (2010) for two sample IV. We derive it in the Appendix and apply it in all subsequent results.

When $\mathbf{Z}$ is a matrix of cohort-year interactions, this approach bears some similarities with cohort averaging methods proposed in Deaton (1985) and Browning, Deaton and Irish (1985). It is worth noting that, even though we plan to include cohort dummies in our instrument matrix in what follows, our approach incorporates several differences. The first is that $\mathbf{X}$ need not be present in the sample that contains data on Y and $\mathbf{Z}$ which will be very useful in estimating leverage effects. The second is that we are not restricted to using discrete grouping variables as instruments: $\mathbf{Z}$ may also contain continous variables.[4] The final

[3]Time dummies or variables interacted with time dummies that are included in $\mathbf{X}$ would also need to be included on the right hand side.

[4]This operationalises a suggestion by Moffitt (1993) in an article discussing how to estimate dynamic models using cross-sectional data.

point is that not all our controls need be included as right hand side variables as they are already 'conditioned out' of the dependent variable. In cohort analysis it is standard to include differences in group averages (e.g. for family size) as right hand side controls. However, as noted in Deaton (1985), these sample averages are noisy measures of the population averages that would ideally be used. This introduces an error-in-variables problem that attenuates estimates of their causal effects which our approach avoids.[5]

4 Empirical analysis

4.1 Data

To investigate the relationship between consumption and leverage, we make use of two 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 5000-8000 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.

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). 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 in all years, while Understanding Society dropped these variables 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 to calculate mortgages in years following 2010. We calculate loan to value ratios by dividing the value of mortgages by the value of homes.

[5] To clarify this point further, imagine that we wished to estimate a model of the form $\Delta Y_t = \beta \mathbf{X}_t + \varepsilon$ and that $\mathbf{X}$ was uncorrelated with ε. The cohort approach takes group averages and then takes first differences. As shown by for example Angrist (1991), this is identical to a 2SLS aproach that uses cohort-year interactions as instruments. It is thus given by $\hat{\beta}_{cohort} = (\hat{\mathbf{X}}'\hat{\mathbf{X}})^{-1}\hat{\mathbf{X}}'\widehat{\Delta Y_t}$. Our estimator would by contrast be equivalent to $\tilde{\beta} = (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\widehat{\Delta Y_t}$.

4.2 Consumption levels regression

We begin with a regression of consumption *levels* on households leverage position. This approach identifies leverage effects by comparing the spending levels across households for given levels of regional house prices. Our specification takes the following form

$$C_i = \gamma_{ct} + \mathbf{X}_i\beta + \alpha_1 \ln(HP)_i + \alpha_2 L_i + \alpha_3(L \times \ln(HP))_i + u_i \qquad (4)$$

where C_i is log consumption, L_i is the household's loan to value ratio, and HP are regional house prices. γ_{ct} are cohort year interactions. We include these to allow for shocks over time that can vary in their effects across age groups. This could include shocks to income expectations which may be correlated with house prices and may be expected to have a larger influence on the spending of the (more leveraged) young than the old. $\mathbf{X}_i$ includes controls for education, sex, region, house type, the number of rooms, the number of children, the number of adults in the household, 4th order polynomials for age and years at current address, and a dummy for having just moved in (moved previous year). We thus control for life-cycle and tenure effects on consumption which we would expect to be strongly correlated with leverage. Including characteristics of house help control for wealth differences between households. The coefficient on regional house prices α_1 can be interpreted as indicating the size of housing wealth effects. However, since our primary focus is on the effects of leverage, this could also be thought of as simply a control that captures regional economic trends. We run this regression for two samples of households: a sample of those with heads aged 25-65, and a subsample of younger heads aged 25-45. In each case we restrict our attention include homeowners who did not move in the current year. Dependent variables are different consumption definitions and employment of the household head.

4.3 Identification issues

As noted above, consumption is observed in the LCF Survey but leverage is not. At the same tim the BHPS includes information on leverage but not on consumer spending. In addition there are reasons to be concerned that leverage and its interaction with house prices (observed only in the BHPS) are endogenous. Initial leverage positions may vary with non-financial assets, unsecured debt and individual home values - all of which could affect consumption responses to house price changes. For these reasons we proceed using the two sample 2SLS approach outlined in section 3.1.

To identify leverage effects seek a source of variation in leverage that explains why some households took out larger loans than others while being not affecting current spending decisions. For this purpose we exploit variation in the average price to income ratios of new loans at the time households moved into their current residences. This variable reflects the cost of credit 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. There is substantial variation in this instrument over time as figure 2 shows.

Figure 3: Credit conditions, 1969-2013

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

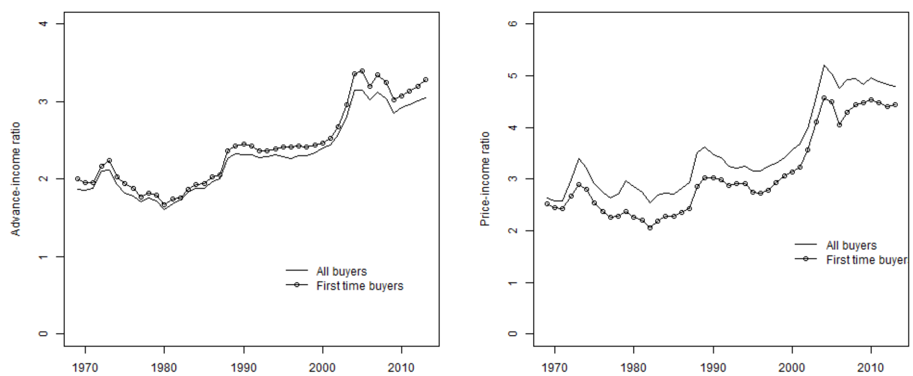

Source: Office for National Statistics

Two requirements must be satisfied in order for our instrument to be considered valid. The first is that it is indeed correlated with the endogenous variables it is replacing, and the second is that the instrument is itself uncorrelated with u_i in (4). The second of these assumptions is normally impossible to verify. Omitted variables are typically omitted because they are unobserved, and so it is impossible to test for an association between them and our instruments (conditional on the exogenous elements of $\mathbf{X}$). However, in the case of two sample IV, it is possible to test for an association between our instruments and a set of characteristics $\mathbf{W}_2$ that may not be observed in sample 1. We do this below. First however we report first stage statistics on the relevance of our instruments.

4.3.1 First stage results

Results for the first stages in our full and younger samples are shown in tables (1) and (2). We have two endogenous variables - leverage and its interaction with house prices. As the discussion in section 3 made clear, it is important that instruments are strong predictors of our endogenous variables. In both first stage regressions, F-statistics are greater than the value of 10 suggested as a rule of thumb by Stock and Yogo (2005). Kleibergen-Paap tests also allow us to reject that the model is underidentified with instruments only predicting variation in one of our endogenous variables.

Table 1: First stage results

	LTV	$LTV \times$ ln(HP)
P/Y (s.e)	0.254	0.566
	(0.064)	(0.329)
$P/Y \times \ln(HP)$ (s.e)	-0.041	-0.068
	(0.012)	(0.060)
Shea partial R^2	0.072	0.070
F-value (p-value)	14.70	15.80
	(0.00)	(0.00)
Kleibergen-Paap (p-value)	28.99	
	(0.00)	
N	32,659	

Table 2: First stage results (25-45)

	LTV	$LTV \times$ ln(HP)
P/Y (s.e)	0.179	0.391
	(0.094)	(0.479)
$P/Y \times \ln(HP)$ (s.e)	-0.020	0.0003
	(0.017)	(0.088)
Shea partial R^2	0.032	0.038
F-value (p-value)	22.92	28.27
	(0.00)	(0.00)
Kleibergen-Paap (p-value)	28.38	
	(0.00)	
N	17,209	

4.3.2 Exogeneity tests

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 rich set of controls for age, tenure in home and cohort-year interactions. Questions about endogeneity may remain however, since some variables are not observed in our second sample. For example, households may have been more likely to move when house prices were high because greater unobservable wealth made them less price sensitive. This would also 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.

To address these concerns we look for an association between our instruments, housing wealth, asset incomes and the probability of being a mover in the BHPS and Understanding Society panels. Table 3 reports results from regressions of these potential sources of endogeneity on our instruments. The instruments are both jointly and individually insignificant in all models suggesting that they are plausibly orthogonal to these omitted variables. Similar results are obtained when we run these regressions on our younger subsample. In addition to these tests we carry our a placebo test of our results on the consumption responses of renters according to the times they moved into their homes below.

Table 3: Exogeneity of instruments

Dependent var.	ln(House value)	Invest inc.>1000	Invest inc.=0	Mover
P/Y	0.08	0.04	0.005	-0.01
	(0.094)	(0.067)	(0.108)	(0.032)
$P/Y \times \ln(HP)$	-0.01	-0.004	-0.003	0.006
	(0.017)	(0.012)	(0.019)	(0.006)
P-values				
P/Y	0.38	0.52	0.97	0.82
$P/Y \times \ln(HP)$	0.44	0.54	0.83	0.92
F-test	0.49	0.79	0.33	0.56
N	37,443	35,170	35,170	27,893
Clusters	11,101	10,604	10,604	8,372

Note: Controls for education, cohort-year dummies, sex, region, house type, number of rooms, number of adults, number of children, a 4th order polynomial in age, a dummy for having just moved in, a 4th order polynomial of years at address, log regional house prices. Standard errors clustered at the individual level.

A further exercise we can do is test to how our instrument compares to leverage positions in the previous period. 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 table (4). While our instrument remains uncorrelated with each of these variables, there is strong evidence that those with higher lagged leverage have greater financial assets and lower home values.

Table 4: Credit conditions vs. LTV_{t-1}

	P/Y	LTV_{t-1}
Mover	-0.01	0.04
	(0.012)	(0.031)
1<Invest inc.<100	0.01	-0.03***
	(0.010)	(0.010)
100<Invest inc.<1000	0.02	-0.06***
	(0.010)	(0.012)
Invest inc.>1000	0.01	-0.09***
	(0.015)	(0.014)
ln(House value)	0.01	-0.05***
	(0.012)	(0.019)
F-test (p-value)	0.60	0.00
R^2	0.74	0.21
N	26,445	19,620
Clusters	8,057	5,834

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: Controls for education, cohort-year dummies, sex, region, house type, number of rooms, number of adults, number of children, a 4th order polynomial in age, a dummy for having just moved in, and a 4th order polynomial of years at address.

4.4 Level regression results

Tables (5) and (6) show results from our level regression for our full and younger samples respectively.

Table 5: Level regression results (full sample)

	$\ln(total + mi)$	$\ln(total)$	$\ln(nondurables)$
$\ln(HP)$	0.227***	0.142***	0.127***
	(0.053)	(0.047)	(0.043)
LTV	-1.734***	-1.076***	-0.827***
	(0.371)	(0.327)	(0.300)
$LTV \times \ln(HP)$	0.445***	0.216**	0.163**
	(0.096)	(0.084)	(0.078)
R^2	0.385	0.358	0.380
N	63,878	63,878	63,878

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: Controls for education, cohort-year dummies, sex, region, house type, number of rooms, number of adults, number of children, a 4th order polynomial in age, a dummy for having just moved in, and a 4th order polynomial of years at address.

Table 6: Level regression results (aged 25-45)

	$\ln(total + mi)$	$\ln(total)$	$\ln(nondurables)$
$\ln(HP)$	0.207***	0.150**	0.130**
	(0.067)	(0.063)	(0.058)
LTV	-1.664***	-1.153**	-0.877*
	(0.537)	(0.506)	(0.464)
$LTV \times \ln(HP)$	0.399***	0.238**	0.171*
	(0.107)	(0.101)	(0.093)
R^2	0.340	0.312	0.327
N	31,394	31,394	31,394

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: Controls for education, cohort-year dummies, sex, region, house type, number of rooms, number of adults, number of children, a 4th order polynomial in age, a dummy for having just moved in, and a 4th order polynomial of years at address.

4.5 Placebo test

We do a placebo test by running a regressing consumption spending of renters on our exogenous variables and instruments. This should pick up any differences between households that become owners in years when credit conditions are loose and those who choose to move into rental accommodation. The instruments are uncorrelated with spending decisions as one would expect.

Table 7: Regressions for renters

$\ln(total - rent)$	Full-sample	25-45
P/Y	-0.02	-0.05
	(0.081)	(0.104)
$P/Y \times \ln(HP)$	0.008	0.012
	(0.015)	(0.019)
F-test (p-value)	0.21	0.51
R^2	0.383	0.331
N	23,331	14,382

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: Controls for education, cohort-year dummies, sex, region, house type, number of rooms, number of adults, number of children, a 4th order polynomial in age, a dummy for having just moved in, and a 4th order polynomial of years at address.

4.6 Differences regression

There may be reasons to think that (4) may be contaminated by unobserved heterogeneity across households that is correlated with our instrument. An alternative model that eliminates fixed effects is the following.

$$\Delta C_{i,t} = \gamma_{ct} + \Delta \mathbf{X}_{i,t}\beta + \alpha_1 \Delta \ln(HP)_{i,t} + \alpha_3 \Delta(L_{t-1} \times \ln(HP))_{i,t} + u_{i,t}$$

We estimate this using three sample instrumental variables as described in section 3.

4.6.1 Controlling for selection

We run estimate our model using households who have not just moved home. This may introduce a selection problem if households who move differ in their leverage positions and spending. To clarify, suppose our model is

$$Y_t - Y_{t-1} = \mathbf{X}_t \beta + u$$

and let the decision to move be determined by the latent variable model

$$y_t^\star = \mathbf{W}_{t-1}\delta + v_{t-1}$$

whereu and v are jointly normal. Individuals move in period t if $y_t^\star > 0$. For identification we want to exploit the relation

$$E[Y_t|\mathbf{Z}=\mathbf{z},y_t^\star<0]-E[Y_{t-1}|\mathbf{Z}=\mathbf{z},y_t^\star<0]=E[X_t|\mathbf{Z}=\mathbf{z},y_t^\star<0]\beta$$

However, we do not observe an estimate of $E[Y_{t-1}|\mathbf{Z}=\mathbf{z},y_t^\star<0]$ in our sample, since we do not know which households in this period will move. We could simply use an estimate of $E[Y_{t-1}|\mathbf{Z}=\mathbf{z}]$ which would ignore possible selection. What we do know however is that

$$E[Y_{t-1}|\mathbf{Z}=\mathbf{z},y_t^\star<0]=E[Y_{t-1}|\mathbf{Z}=z]+\rho E[v_{t-1}|\mathbf{Z}=\mathbf{z},y_t^\star<0]$$

$$E[Y_{t-1}|\mathbf{Z}=\mathbf{z},y_t^\star<0]=E[Y_{t-1}|\mathbf{Z}=\mathbf{z}]+\rho E[v_{t-1}|\mathbf{Z}=\mathbf{z},v_{t-1}<\mathbf{W}_{t-1}\delta]$$

$$E[Y_{t-1}|\mathbf{Z}=\mathbf{z},y_t^\star<0]=E[Y_{t-1}|\mathbf{Z}=\mathbf{z}]+\rho E[\lambda(\mathbf{W}_{t-1}\delta)|\mathbf{Z}=\mathbf{z}]$$

where $\lambda(.)$ is the inverse mills ratio. Selection means our regression omits the term $\rho E[\lambda(\mathbf{W}_{t-1}\delta)|\mathbf{Z}=\mathbf{z}]$ from the left-hand side so we include it as a the right-hand side variable to control for selection. We include in $\mathbf{W}_{t-1}$ a dummy for whether the households has children of school age, along with our instrument set. We theorise that households with children of school age will be less likely to move in response to economic shocks, because it can disrupt their education. $\mathbf{W}_{t-1}\delta$ can only be estimated using the BHPS as it requires panel data. The school age dummy enters highly signficantly.

4.7 Three sample results

[TO BE FILLED IN]

Table 8: Three sample results (full sample)

	$\ln(total+mi)$	$\ln(total)$	$\ln(nondurables)$
$\Delta\ln(HP)$			
$\Delta(\ln(HP)\times LTV_{t-1})$			
R^2			
N			
* $p<0.10$, ** $p<0.05$, *** $p<0.01$			

Note:

5 Discussion

[TO BE FILLED IN]

References

[1] References