



BANK OF ENGLAND

Working Paper No. <xxx>

The role of debt in UK household spending decisions

Philip Bunn⁽¹⁾ and May Rostom⁽²⁾

PRELIMINARY AND INCOMPLETE – PLEASE DO NOT QUOTE

Abstract

Household debt rose sharply in the United Kingdom over the decade before the financial crisis. This paper uses household level microdata to investigate how debt affected consumption before and after the crisis. Whilst the increases in debt were largely matched by an accumulation of financial assets, we find evidence that debt may have provided some support to consumption over the period when debt rose rapidly. But since 2007, the extent to which high debt households consume more has fallen back, which will have weighed on consumption growth over this period. Weaker consumption among high debt households since 2007 may be related to debt amplifying the impact of income shocks.

Key words: <optional>

JEL classification:

(1) Bank of England. Email: philip.bunn@bankofengland.co.uk

(2) Bank of England. Email: may.rostom@bankofengland.co.uk

The views expressed in this paper are those of the authors, and not necessarily those of the Bank of England. We are grateful to <acknowledgements> for useful comments. This paper was finalised on <Press Office clearance date>.

The Bank of England's working paper series is externally refereed.

Information on the Bank's working paper series can be found at www.bankofengland.co.uk/publications/workingpapers/index.html

Publications Group, Bank of England, Threadneedle Street, London, EC2R 8AH
Telephone +44 (0)20 7601 4030, Fax +44 (0)20 7601 3298, email mapublications@bankofengland.co.uk.

Contents

Summary	3
1 Introduction	5
2 Theory and literature	6
3 An overview of developments in household balance sheets	Error! Boo
3.1 Trends on housheold balance sheets	Error! Boo
3.2 How have developments in household balance sheets affected consumption?	Error! Boo
4 Research design and data	Error! Boo
4.1 Research design	Error! Boo
4.2 Data	Error! Boo
5 Results	Error! Boo
5.1 Econometric results	Error! Boo
5.2 Estimating the impact of debt on aggregate consumption	Error! Boo
6 Understanding how debt affects consumption	Error! Boo
6.1 Causality versus differing expectations	Error! Boo
6.2 Evidence on why households have cut spending since 2007	Error! Boo
7 Conclusion	Error! Boo
Appendix	11
References	12

Summary

<Summary text – maximum length two pages>

1 Introduction

A major development in household balance sheets over the period between 1992 and 2007 was the build up of household debt. The stock of debt more trebled from around £450bn in 1992 to around £1.5 trillion in 2007 with most of the increase coming from the late 1990s onwards. In relation to income, debt rose by around 60% over this period.

Debt is not normally part of most mainstream macroeconomic models and the extent to which debt may have implications for aggregate spending is not well understood. In a standard model life-cycle type model debt has no influence on households' spending decisions: households typically borrow to smooth through their consumption, which is determined entirely by their expected lifetime income and wealth. But debt could potentially affect aggregate consumption, at least temporarily, if households' ability to borrow to bring forward their consumption changes. Or debt may affect how aggregate consumption responds to shocks if indebted households respond to those shocks in a different way to other households.

In the context of the UK economy, it is possible that the large build up of debt over the period leading up to the financial crisis helped to finance household consumption, and that the recession was deeper, and the subsequent recovery slower, as a result of indebted households making larger adjustments to spending than other households. Our work aims to investigate these issues. Evidence suggests that high debt levels did depress consumer spending in the US after 2007 (Dynan (2012) and Mian, Rao and Sufi (2013)).

It is difficult to assess how debt has affected consumption using aggregate data and therefore it is important to use microdata to look at how households with different levels of debt have adjusted their consumption. We use microdata from the Living Costs and Food Survey to investigate how the consumption of high and low debt households has differed over the period since 1992. This approach cannot prove whether debt directly causes consumption or not, only whether high debt households consume more or less. But we also use evidence from the Bank of England/NMG survey to assess whether debt has led some households to reduce their spending since the financial crisis.

We find that households with high debt levels have, on average, consumed more than those with low debt. Whilst increases in debt in the decade prior to the financial crisis were largely matched by an accumulation of financial assets, debt may have provided some support to consumption over the period when debt rose rapidly. But since 2007, the extent to which high debt households consume more has fallen back, which will have weighed on consumption growth over this period. We find evidence that weaker consumption among high debt households since 2007 may be related to debt amplifying the impact of income shocks.

The paper is organised as follows. First we discuss the theory and previous literature around the question of whether debt does affect consumption. We then provide an overview of developments in household balance sheets in the UK over the period since 1992. We then move on to discuss our research design, the data and then the econometric results on how debt effects



on. Then we discuss the interpretation of our results and bring in some evidence on why households may have cut spending after 2007 from the NMG survey. Finally, we conclude.

2 Theory and literature

In the permanent income/life-cycle model households borrow or save in order to smooth their consumption over their life cycle. Current consumption depends on expected lifetime income and wealth. Households with debt do not respond differently to those without, whether a household has debt is simply related to how their current income and wealth relates to their average income and wealth over their lifetime. In the simplest form of the life-cycle model, households can borrow as much as they choose, the cost of borrowing is constant, households are assumed to be able to accurately predict their lifetime income, there is no uncertainty and expectations do not change.

In practice, many of the assumptions that underlie the life cycle model may not hold. For example, households are not certain about their lifetime incomes; they may have unrealistic expectations or may be affected by shocks which change those expectations. They may also not be able to borrow against their future income, or be able to borrow as much as they would like to be able to fully smooth their consumption, and their ability to borrow or the cost of borrowing may change over time. These factors could, in principal, help to explain why indebted households could respond differently to shocks to things like interest rates, income or employment.

There are reasons to think that indebted households may have responded differently to the shocks associated with the financial crisis. Benito et al (2007) conclude that there was little evidence to suggest that high and low debt households responded differently to changes in their financial positions over the period from 1997 to 2004. But that was over a period of relative stability and the shocks experienced by households since 2007 are likely to have been much larger. For many households, the financial crisis is likely to have lowered expected income or made it more uncertain. Highly indebted households are more vulnerable to these types of shocks because their debt servicing costs are likely to be a higher proportion of their income. Households may become more concerned about their ability to service their debts, and therefore choose to reduce consumption by more than would be implied by the income shocks alone, in order to either reduce their indebtedness or build up a bigger buffer of savings to meet future repayments.

The literature on how debt might affect spending does back to Fisher's (1933) debt deflation theory, which says that the US Great Depression was caused by falling prices increasing the real burden of debt, which led to further deflation. This explains how debt helped to amplify the initial shock as it propagated through the economy. King (1994) explains how Fisher's debt deflation theory might have been relevant in explaining weakness of consumption following 1990s recession. He puts forward a model of precautionary saving in which agents who had borrowed on the expectations of future returns suffer adverse shocks that lead them to consume less and repay debt. Other agents experience offsetting positive shocks, but they do not increase

consumption by enough to fully offset the effect on aggregate, which implies debt levels can weigh on aggregate consumption.

Since the recent financial crisis, new research on the role of debt has typically tried to incorporate some type of reduction in households' ability to borrow. For example, Krugman and Eggertson (2012) assume that there is a limit on how much debt individuals can hold, and if that limit is revised down, highly indebted households are forced to reduce spending sharply with no offsetting response from non-debtors. Other models such those of Guerriero and Lorenzio (2011) and Philippon and Midrigan (2011) also try to explain weak consumption by incorporating a tightening in the ability of households to borrow to smooth consumption and by a reduction the ability of households to withdraw equity from their homes.

Empirical work in the US has found evidence of an empirical relationship between high pre-crisis debt levels and weak post crisis consumption. Mian, Rao and Sufi (2013) show that the decline on US consumption following the crisis was greater in counties with higher leverage prior to the crisis. Dynan (2012) uses household level data to show that highly leveraged households experienced larger declines in spending between 2007 and 2009, after controlling for wealth effects and other factors.

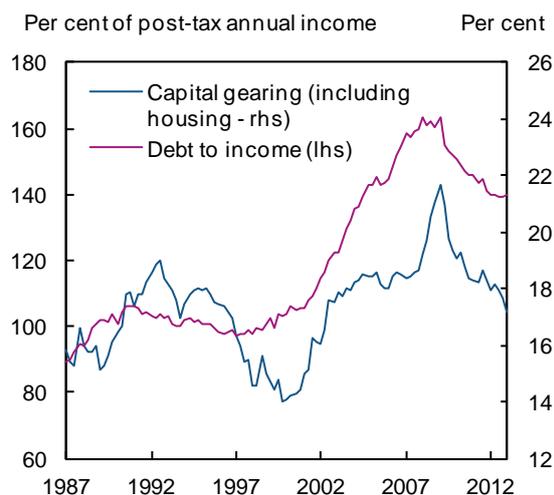
But there is no such empirical work for the UK looking at whether high debt levels have weighed on consumption since the crisis. Our work aims to try and fill that gap, and also say something about the extent to which debt supported consumption during the period before the crisis.

3 An overview of developments in household balance sheets

Over the period from 1992 to 2007, household debt rose from around 100% of annual income to a peak of around 160%, although the debt to income ratio only began rising from around 1999 onwards (Chart 1). Those increases in debt were also accompanied by an accumulation of financial assets and increases in asset values, such that measures of capital gearing in 2007 were close to their 1992 level, although debt did rise faster than asset holdings in the 2000s having grown more slowly during the 1990s (Chart 1).

Since 2007, the stock of household debt has stabilised, and the debt to income ratio has fallen back from its peak, given continued, albeit modest, growth in nominal incomes. Measures of capital gearing rose during the crisis as asset values fell sharply, but as asset prices recovered, those measures of debt relative to asset values fell back.

Chart 1: Household debt to income and capital gearing



This section of the paper provides a more detailed discussion of these developments in household balance sheets and then considers how they may have affected household spending.

3.1 Trends in household balance sheets

Mortgage debt accounts for around 80% of total household debt, and increases in the aggregate debt to income ratio since the late 1990s largely reflect higher levels of secured debt (Chart 2). Annual growth in secured net lending rose from around 5% in 1993 to a peak of around 15% in 2004 and remained above 10% until the start of the financial crisis in 2008 (Chart 3). Unsecured lending grew more rapidly than secured lending during the 1990s. Over the 2000s prior to the crisis, secured and unsecured lending increased at a broadly similar rate, although growth in unsecured debt did start to fall back from around 2005. Since 2009, growth in both secured and unsecured net lending has been close to zero.

Chart 2: Household debt to income ratio

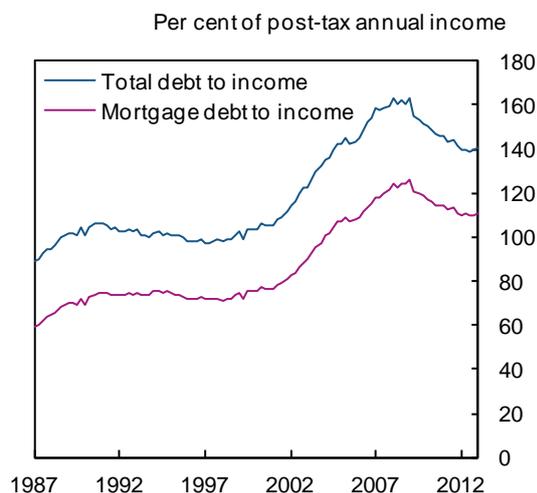
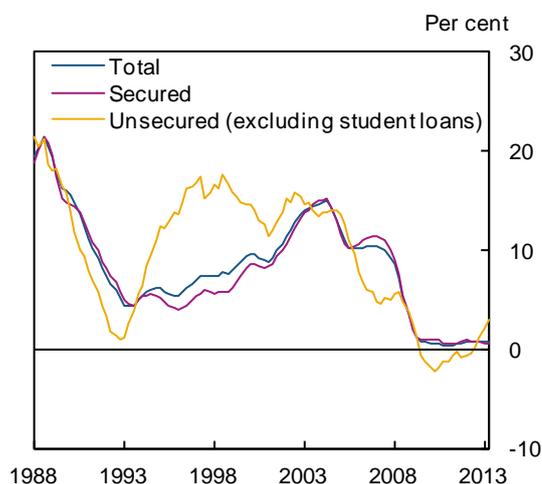
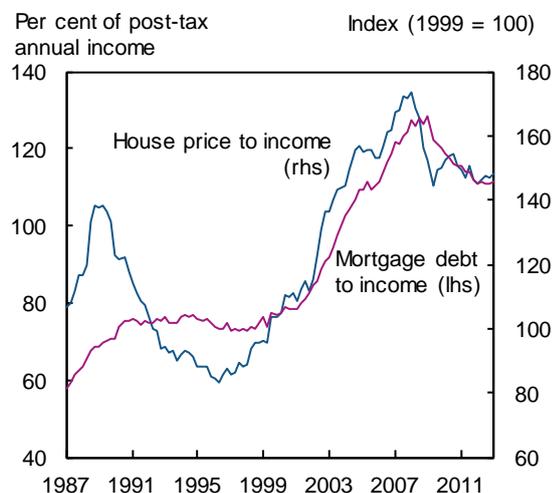


Chart 3: Annual growth in net lending to households



The increase in household debt to income from the late 1990s appears to have been largely related to increases in house prices. That is apparent from the fact that the mortgage debt to income ratio has moved closely with the ratio of house prices to income since 1992 (Chart 4). Higher house prices mean that new entrants into the housing market have to pay more to buy their house. If sellers were using the proceeds to pay down their debt there would be no overall effect on aggregate debt. But as last time sellers tend to have either small mortgages or no mortgage, higher house prices tend to lead

Chart 4: House prices to income and mortgage debt to income



to higher secured debt.

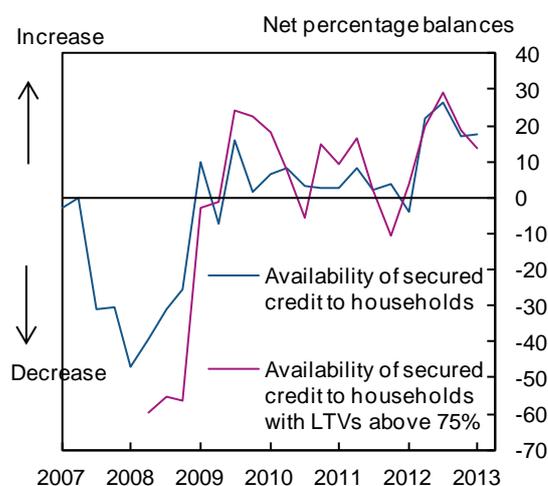
Although higher house prices have led households to take on more debt since the early 1990s, factors that have increased households' ability to borrow may also have themselves contributed to increases in house prices. For example, lower real interest rates will have made it cheaper for households to borrow. There was a substantial and sustained fall in long-run risk-free real interest rates in the late 1990s (Chart 5). Waldron and Zampolli (2010) argue that the fall in real interest rates is a key factor behind the increases in both house prices and debt from the late 1990s.

On top of lower real interest rates, looser credit conditions made it easier for households to borrow in the period before the crisis. For example, banks became more willing to lend to households at high loan to income ratios. In part, that could also reflect lower and more stable inflation, which when combined with lower nominal interest rates removes the problem of front end loading where initial mortgage repayments are unaffordable at high loan to income ratios that is present with high inflation and nominal interest rates.

Chart 5: Five year real interest rates, five years forward



Chart 6: Credit availability from the Credit Conditions Survey



Since 2007, credit conditions have tightened substantially, and that is one factor that is likely to explain why the stock of mortgage debt has stabilised since then. Evidence from the Bank's Credit Conditions Survey suggests that credit availability fell significantly, particularly at higher loan to value ratios (Chart 6), and credit spreads also rose. The resulting fall in the number of first-time buyers entering the housing market will have weighed on the flow of new lending. And to the extent that tighter credit conditions have contributed to falls in house prices during the crisis, that will also have reduced the size of mortgages taken on by the first-time buyers who remained.

In aggregate, the build up in household debt over the period from 1992 to 2007 was largely matched by a build up in assets. As house prices rose, households needed to take out bigger

mortgages to be able to purchase a house. But financial assets were acquired at a similar rate to liabilities, which suggests that the proceeds from those sales were largely left on deposit or invested in financial assets (Chart 7).

Chart 7: Household net acquisition of financial assets and liabilities **Chart 8: Household assets and liabilities**

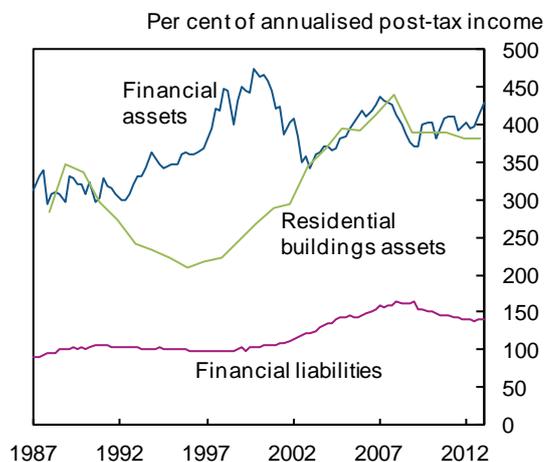
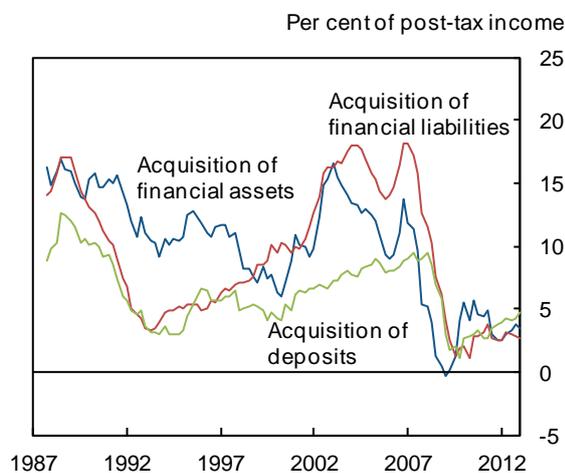
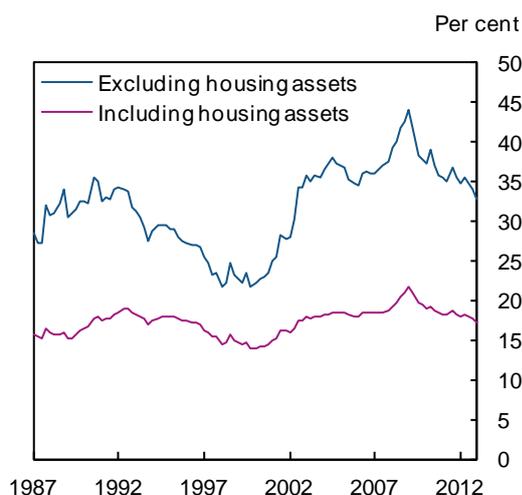


Chart 9: Household capital gearing



The overall balance sheet position of households also depends on changes in asset values as well as patterns of asset and liability accumulation. Changes in equity prices can have a significant impact on gross financial wealth. For example, financial wealth increased significantly during the second half of the 1990s as equity prices rose strongly, but it fell back again during the dot com crash as prices fell back (Chart 8). Movements in house prices also affect the value of housing assets, and the value of those assets increased significantly between the late 1990s and 2007 as house prices rose. Taking together changes in debt and asset holdings, measures of capital gearing - which measure the stock of debt in relation to the value of assets - were at a broadly similar level in 2007 to in 1992, even if housing assets are taken into account, although those 1992 levels were themselves elevated (Chart 9). Within this period gearing fell back during the 1990s and rose during the 2000s as debt levels increased. But overall, these measures have remained below one, indicating that households have always had significant net wealth.

Although higher debt levels have not been associated with a significant deterioration in the overall state of household balance sheets, they have led to notable changes in the distribution of balance sheets. Increases in borrowing by one part of the population were largely matched by a big increase in saving by another part of the population. This is clear from looking at changes in debt and wealth between 1995 and 2005 from the British Household Panel Survey microdata. Those trading up in the housing market, who tend to be younger households, were made worse off by the increase in house prices because they had to pay more for their houses and they accumulated more debt (Chart 9). At the other end of the spectrum, those who either already owned large houses or who traded down and sold a higher values, who are more likely to be older households, saw significant increases in their wealth (Chart 10).

Chart 9: Household debt by age

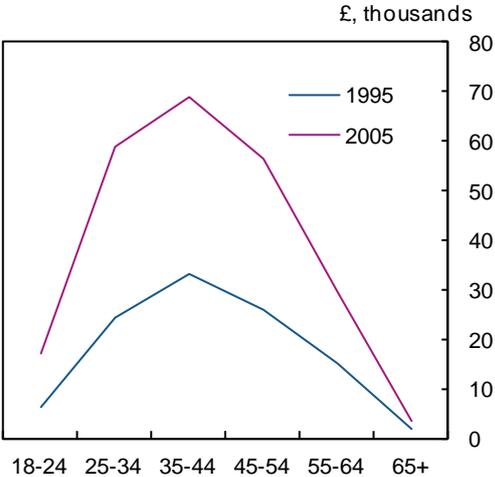
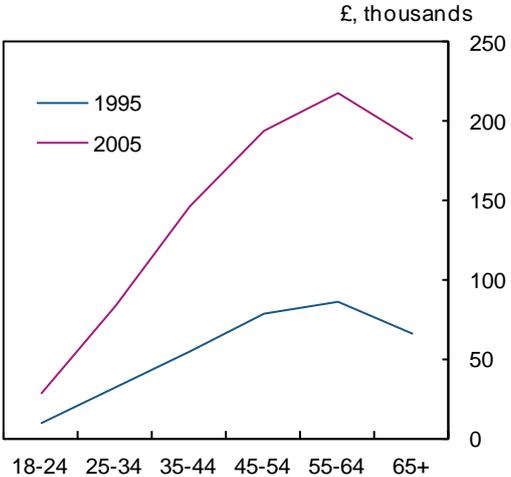


Chart 10: Household gross financial wealth plus housing assets by age



3.2 *How have developments in household balance sheets affected consumption?*

It is not apparent from aggregate data that the build up of household debt did significantly boost consumption in the period before the crisis. At a very simple level, one way to look at whether debt may have boosted consumption is to compare the 1992-1998 period when debt was growing broadly in line with income to 1999-2007 when debt rose sharply relative to income. But average consumption growth is very similar over the two periods at around 3.5%, although both are a little above the long run average of 2.8%. It is clear from this that there was no consumption boom like in the late 1980s.

Household spending depends on both income and the share of income saved, and if debt were being used to finance consumption that should be reflected in a falling saving ratio. The household saving ratio did fall significantly, by around 8 percentage points, over the 1992-2007 period, but the rate of decline was similar between 1992 and 1998 when debt to income was not rising to between 1999 and 2007 when it was increasing. Again that does not suggest that there was a debt fuelled consumption boom. But this analysis is very simplistic and there are lots of other factors that affect consumption growth - it does not say debt has had not effect, but it

probably does rule out debt being the most important factor driving trends in aggregate consumption.

Chart 11: Consumption growth

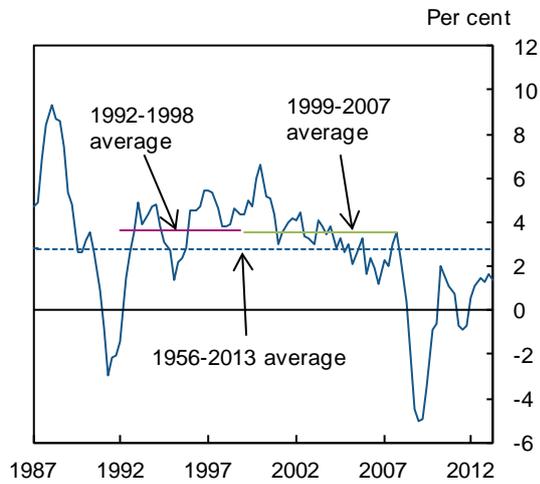
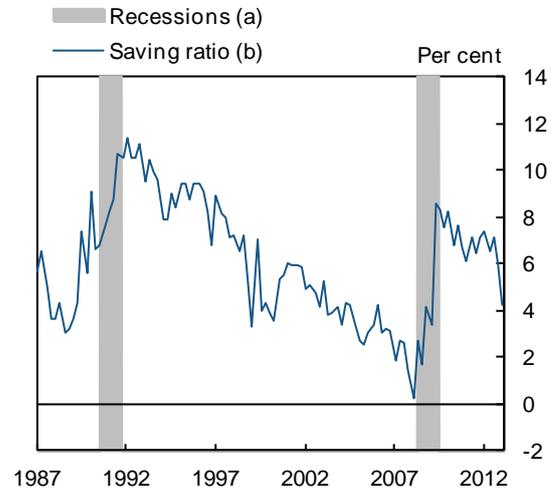


Chart 12: Household saving ratio



As increases in debt were largely matched by an accumulation of assets it may be more useful to consider the relationship between consumption and debt that is actually available for consumption. Housing equity withdrawal occurs, whenever households, in aggregate, increase borrowing on secured assets without spending the proceeds on improving or enlarging the housing stock, and is an indicator of the debt that is available for consumption. There is a correlation between housing equity withdrawal and consumption which might imply some of this borrowing was used to finance consumption (Chart 13). Unsecured borrowing can also be used to finance consumption, and that also picked up over a period when consumption was rising relative to income.

Chart 13: Housing equity withdrawal and consumption

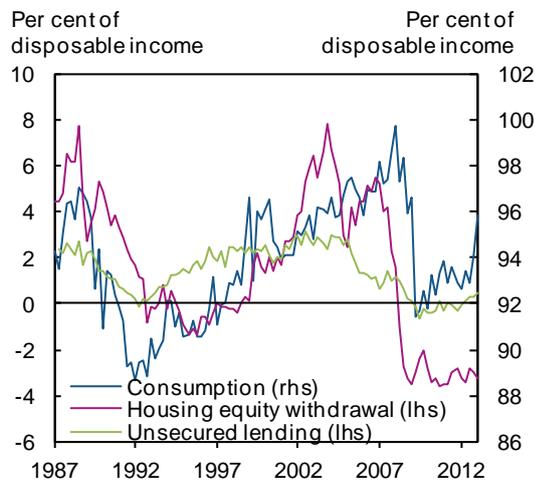
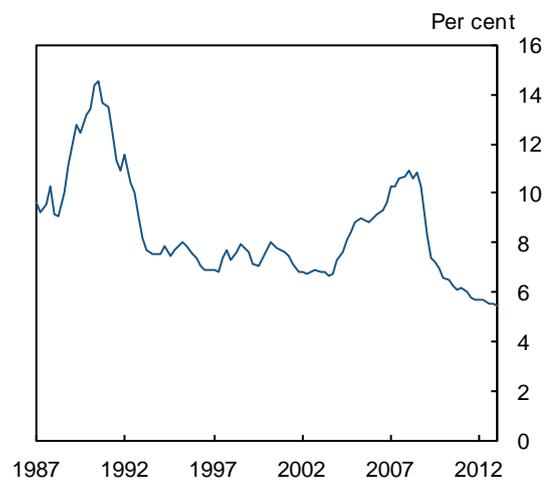


Chart 14: Income gearing



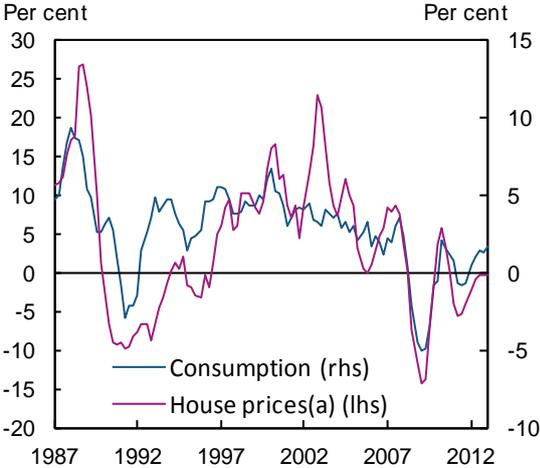
The cost of servicing debt may also play a role in consumption decisions. Higher debt servicing costs arising from high debt levels may lead to lower consumption because they mean that households have less income available to spend on goods and services. Income gearing – the cost of servicing debt relative to income – was low between 1992 and the early 2000s, but it did start to pick up from around 2003 as monetary policy tightened (Chart 14). Increases in interest payments are likely to have weighed on consumption over this period. Since the crisis, debt servicing costs have fallen to historically low levels on the back of loose monetary policy, which will have helped to support spending.

Changes in the asset side of the household balance sheet as well as changes in liabilities may have implications for household spending. Changes in financial asset prices, primarily equity prices, may lead households to reassess their future wealth and adjust spending in response. Equity prices rose during the 1990s and may have supported consumption over that period.

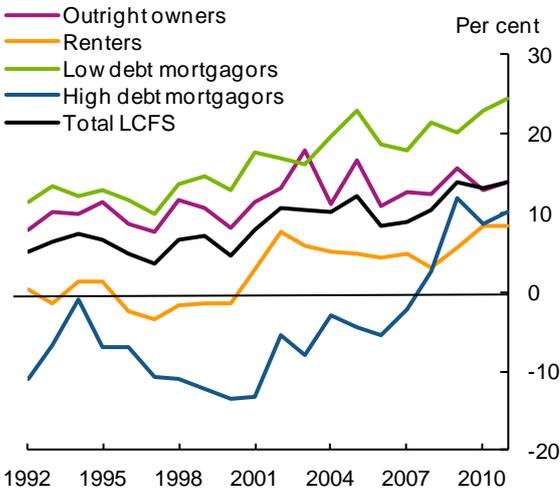
Changes in house prices can affect spending in a number of ways, including by affecting the value of collateral which households can borrow against to fund consumption. Consumption could also be affected if different types of people respond in different ways to changes in house prices. For example, older households who are likely to trade down are made better off by increases in house prices and may increase their consumption in response. But whether there is an impact on aggregate spending also depends on whether those households who want to get into the housing market or trade up reduce their spending by a similar amount. In the data, house prices and consumption are well correlated (Chart 15) and increases in house prices could have supported consumption in the period before the crisis. But there is no consensus in the literature on whether that reflects causality (as argued by Campbell and Cocco (2007)) or whether the two simply move together because they are determined by similar factors such as expectations of future income (Attanasio et al (2009)).

Chart 15: Real house prices and consumption

Chart 16: Saving ratios of different groups of households



(a) House prices series is the ONS house price index deflated by the consumption deflator.



Overall, it is difficult to assess how debt has affected consumption using aggregate data. It is important to use microdata to investigate how consumption varies across households with different levels of debt. Chart 16 shows saving ratios for different groups of households calculated from the Living Costs and Food Survey (see section 4.2 for further details on this survey). According to the LCFS, high debt mortgagors (defined as those with a debt to income ratio 2) had negative saving ratios throughout the period from 1992 to 2007, ie their consumption was on average higher than their income. That would be consistent with these households using debt to finance some of their consumption. Since 2007, the saving ratio for households with high mortgage debt has risen by more than for other households, which could be consistent with a larger response by these households to the shocks associated with the financial crisis. The research presented in the rest of this paper looks at these issues in more detail.

4 Research design and data

4.1 Research design

We investigate how debt affects non-housing consumption by estimating household level consumption equations that incorporate a mortgage debt to income variable. Our econometric approach allows us to test whether debt has had an impact on consumption after controlling for all observable characteristics.¹ The coefficient on the leverage variable is allowed to vary over time to investigate whether debt has had a different impact before and after the crisis.

Our framework is based on a life cycle type model that is widely used in the housing and consumption literature, except that we adapt it to also include leverage. For example, Attanasio et al (2009) use a similar methodology to examine the relationship between housing and consumption. As discussed above, although debt does not play an important role in the standard life cycle model there are reasons to think that some of those assumptions may not hold, and therefore it is useful to include debt in the model to test whether that holds.

We estimate the following equation at the household level:

$$c_{it} = \beta_1 leverage_{it} + \beta_2 leverage_{it} * year_{it} + \beta_3 year_{it} + \beta_4 cohort_{it} + \beta_5 X_{it} + e_{it}$$

Where C_{it} is the log of real non-housing consumption for household i at time t , *leverage* is the ratio of outstanding mortgage debt to households disposable income. *year* is a vector of time dummies to capture time specific shocks, and *cohort* is vector of 5-year date-of-birth buckets. X is a vector of controls for income, region, age (up to the 5th power), number of children and adults in the household, tenure, ethnicity, marital status, gender, occupation, school leaving age, and regional house prices.

The β_1 coefficient measures the common effect of leverage on consumption for all time periods while the vector β_2 captures how that impact changes over time. As our data is a repeated cross-

¹ By controlling for as many factors as possible, we are aiming to compare the consumption of two households of similar income, age, education etc, but where one household has high debt and the other has low debt.

section, we are unable to control for unobserved individual heterogeneity using standard panel data methods. But we do include cohort dummies (based on 5 year birth intervals) which capture a pseudo fixed effect for all the households within a cohort.

Our approach allows us to test whether high debt households consume more or less than low debt households at different points in time, but it does not allow us to conclusively say that debt directly caused high debt households to consume more or less. That is because debt and consumption are potentially endogenous. Households with high debt may use debt to finance consumption or may respond differently to shocks, which would imply a link. But alternatively, households with more optimistic expectations of future income are may be more willing to take on debt and consume more because of those expectations.

To try and address this issue about whether there is a causal link between debt and consumption, we make use of the Bank of England/NMG survey. The survey asks households a range of questions about their finances and can be used to investigate whether suffering income shocks makes household with debt more likely to cut spending or whether high debt households are just more likely to have received lower income than they expected.

4.2 Data

We use household level microdata from the Living Costs & Food Survey (LCFS) between 1992 and 2011, previously known as the Family Expenditure Survey, to investigate how debt affects consumption. This is an annual, repeated cross section survey, rather than a pure panel and it contains information on consumption, income, mortgage debt and a range of other household and individual characteristics. Data are available from 1961-2011, but we only use the last twenty years in our analysis as mortgage debt data is only available from 1992. There are some data on unsecured debt, but they are not comprehensive and so we focus only on secured debt.

On average, the LCFS contains approximately 5,300 households per year, although the sample size does decline gradually over time. The survey also switches between calendar and fiscal year collection more than once. For consistency, we convert the data to a calendar year setup. We only include households where the head of the household is aged between 21 and 70 and was born between 1926 and 1985.

Our analysis uses weekly non-housing consumption. Consumption data in the LCFS is collected from a detailed diary that households keep for two-weeks. We specifically use consumption expenditure net of housing because the way housing consumption is measured in the LCFS is not consistent with the methodology used in the National Accounts – for homeowners, the LCFS only measures mortgage payments rather than using a measure of imputed rents like the National Accounts.

We use the 2012 Bank of England/NMG survey to investigate whether there is a causal link between debt and consumption. This is a cross-sectional survey of around 4000 households commissioned by the Bank. It does not contain consumption data, but it does include information on debt, income, whether a household has cut spending because of concerns about

their debt and whether income has turned out lower than expected. Whilst the NMG survey has been running since 2004, the questions we use in this analysis have not all been asked in earlier years.

5 Results

5.1 *Econometric results*

On average, households with high secured debt levels have higher non-housing consumption than those with low debt. In our baseline specification (equation (1) in Table A in the appendix), a household with a mortgage debt to income of 3 rather than 2 can be expected to consume around 2% more, all else equal. Outright owners consume more than mortgagors, whilst renters consume less.

We find that the impact of debt on consumption does vary over time. Since the 2007, the degree to which households with high debt levels have consumed more has fallen back, although households with high debt do still consume more. The smaller boost to consumption from high debt households after the crisis will have weighed on consumption growth over this period. Chart 17 summarises the coefficients on the year/leverage interaction variables which show how the impact of a one unit change in debt to income (such as from 2 to 3) on consumption has changed relative to 2007. The impact was around 1.5pp smaller after 2007 than it was in 2007 and those differences are statistically significant (except for in 2010). Precise coefficients are reported in Table A in the appendix.

The impact of debt on consumption following the crisis is estimated to be similar to that following the early 1990s recession. The coefficients on the year/leverage interaction variables are comparable between 2008 and 2011 to what they were between 1992 and 1998. But the effect of debt on consumption in the early 2000s was significantly larger, which might be consistent with debt supporting consumption over this period. Taking together the common and the time varying effects, a household with a debt to income ratio of 3 rather than 2 is estimated to have consumed about 1% more during the 1992-1998 period, which rose to around 4% in the early 2000s, before falling back to around 1.5% after the crisis (Chart 18).

Chart 17: Impact of 1 unit increase in debt to income ratio on non-housing consumption, relative to 2007

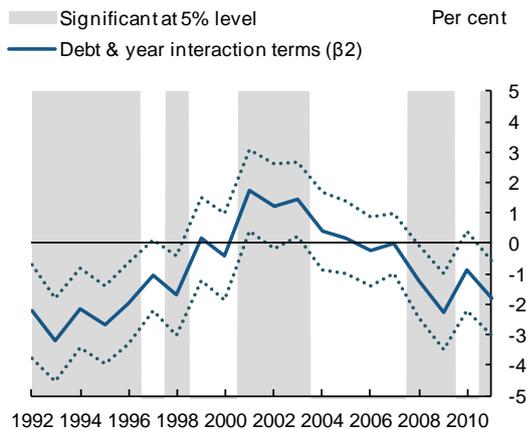
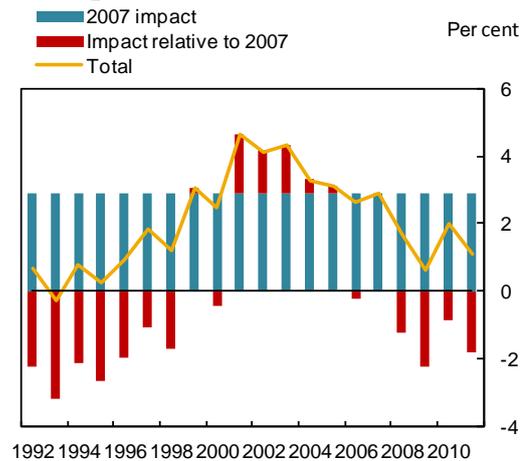


Chart 18: Impact of 1 unit increase in debt to income ratio on non-housing consumption



The larger bigger of debt on consumption in the early 2000s largely reflects higher spending on durables (Chart 19). That spending could be associated with borrowing to fund spending on housing related items on furniture. But relative to 2007, the impact of debt on both durable and non-durable consumption has fallen by a broadly similar amount (Charts 19 and 20), although the durables estimates are not generally statistically significant.

Chart 19: Impact of 1 unit increase in debt to income ratio on durable consumption, relative to 2007

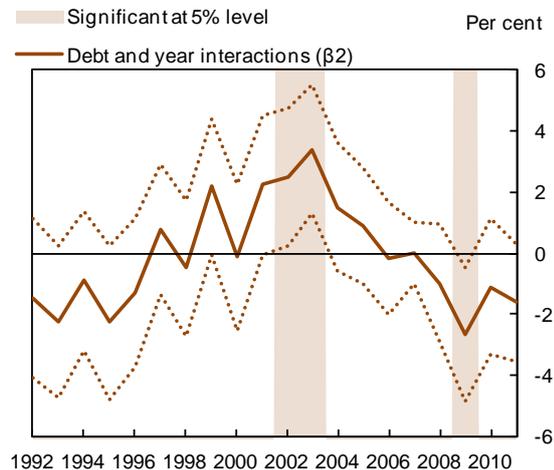
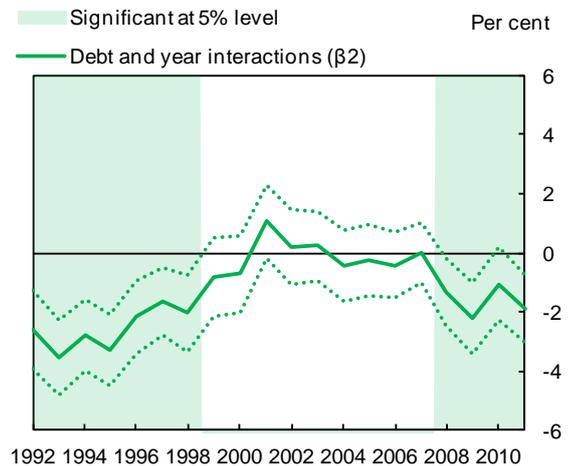


Chart 20: Impact of 1 unit increase in debt to income ratio on non-durable consumption, relative to 2007



Our results suggest that debt may have had a smaller impact on consumption in the UK during the crisis than in the US. The paper most comparable to our work is Dynan (2012). She finds that a household with a debt to income ratio of 3 rather than 2 had 6% lower consumption growth between 2007 and 2009. That is around two and half times larger than our estimate of 2.3%. Dynan (2012) only looks at changes in consumption over the crisis period and so does not examine whether high debt households always have a higher level of consumption.

It is also important to be clear how our results relate to the concept of deleveraging. Deleveraging may be thought of as households actively paying down debt over and above their regular repayments. In aggregate, the stock of mortgage debt has been broadly flat since 2008 and there is little evidence of an increase in overpayments. Our work says nothing about the implications of making overpayments of debt, we are simply estimating the impact of having debt on consumption. The smaller impact of debt on consumption post 2007 could be consistent with some highly indebted households repaying debt, but equally they could just be building up their savings balances.

5.2 *Estimating the impact of debt on aggregate consumption*

Our regression results can be used to estimate the impact of how mortgage debt affect has affected aggregate consumption. We estimate this by comparing the actual fitted values from our equation to fitted values where we do not allow debt to affect consumption, ie we set the coefficient on that variable and the leverage/year interactions to zero. These estimates are shown in Chart 21, which is scaled up from the LCFS to match non-housing consumption in the National Accounts.

Our results suggest that debt had little impact on consumption for most of the 1990s. But between 1998 and 2007, as both the extent to which high debt households consumed more increased and debt levels rose, our estimates suggest that debt did support consumption. There is some volatility in the estimate from year to year, but our results are consistent with debt adding around 3% to the level of non-consumption over this period, which would equate to around 0.3% a year on consumption growth on average. That compares to annual average consumption growth of around 3.5% over this periods (Chart 11). Mapping the results into saving ratio space (and taking housing and NPISH consumption into account) they imply increased spending by high debt households explains virtually none of the fall in the saving ratio between 1992 and 1998, but it can potentially account for around half of the fall over the period between 1998 and 2007 (Chart 22). So whilst the increases in debt were largely matched by an accumulation of financial assets, there is some evidence debt may have provided modest support to consumption over the period when debt rose rapidly.

Chart 21: Estimated impact of debt on non-housing consumption

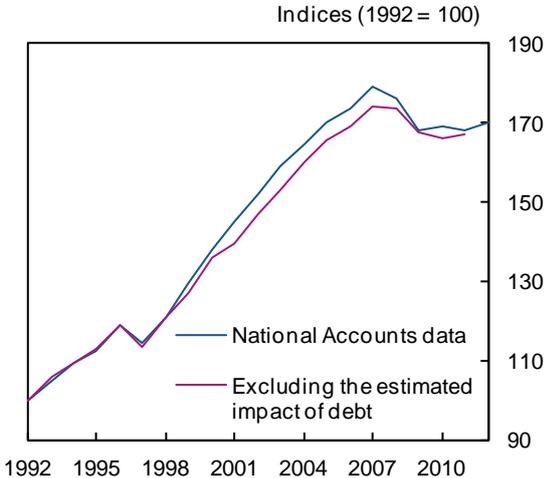
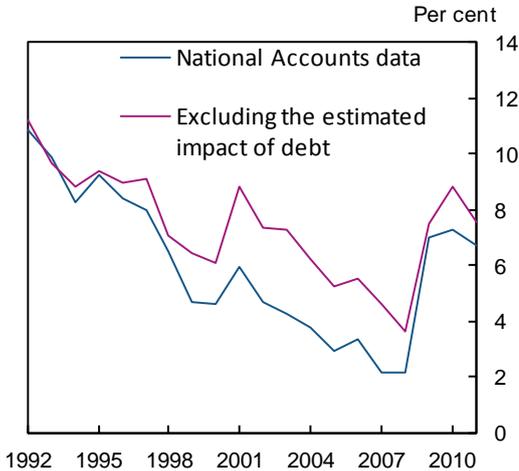


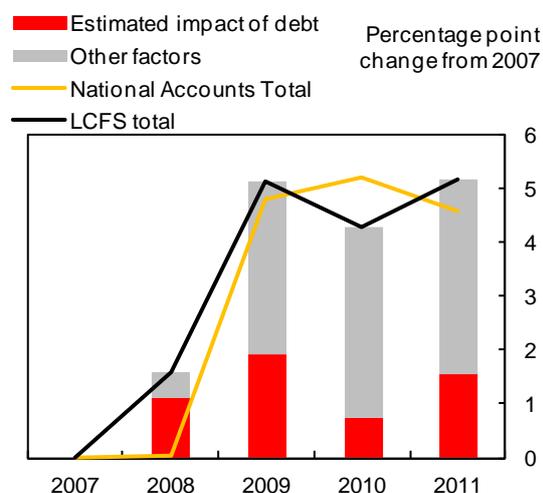
Chart 22: Estimated impact of debt on the saving ratio



Since 2007, the extent to which debt has supported consumption has fallen back, and is estimated to have reduced aggregate non-consumption by around 2%. In saving ratio space we estimate that changes in the impact of debt added around 1.5pp to the aggregate household saving ratio, out of a total increase of around 5pp (Chart 23). Whilst that is a significant contribution to the rise in the saving ratio it still leaves a considerable part of the increase to be accounted for by other factors.

These estimates of how debt has affected consumption could be an underestimate of the impact of debt because our analysis has focussed solely on mortgage debt, and it is possible that the effects would be larger if unsecured debt could be taken into account. The other important caveat to this is that for some households to hold debt, others have to own assets, and these estimates do not take into account any response by non-debtors.

Chart 23: Contributions to the change in the saving ratio from 2007



6 Understanding how debt affects consumption

6.1 Causality versus differing expectations

The finding that high debt households, on average, consume more than low debt households could simply reflect those households borrowing more in order to consume more (as well as buy a house). But debt does ultimately need to be repaid, and therefore it is unsustainable that high debt households consistently use debt to finance their consumption. It may be that debt has been used to finance unsustainably high consumption over our sample period. But it could also be the case that households use mortgage debt to bring forward their consumption shortly after buying a house when their leverage is relatively high and they do not use debt to finance consumption later on when they are close to repaying the loan and their debt to income ratio is lower.

Our results do not prove that there is a causal link between debt and consumption that – they are only consistent with the hypothesis. There are also other possible explanations for why high debt households could spend more than low debt households that could be consistent with debt not having a causal relationship with consumption. These include:

1) High debt households expect higher future income growth. If those who take on high levels of debt expect higher future income growth than those with lower debt, they have a higher expected lifetime income and so should also have higher current consumption, even if their current incomes are the same. Back of the envelope calculations suggest this is plausible

because the required differences in expected future income growth required to match our results are small.²

2) High debt households expect higher future house price growth. If high debt mortgagors expect stronger future house price growth than households with lower debt, they should expect to make larger capital gains, and it would be rational for them to consume some of the expected gains. Again, the required differences in expectations are relatively small, although the capital gains do need to exceed the cost of borrowing for there to be a surplus available to consume.³ But unlike with income expectations, actual house price growth should not vary significantly across high and low debt households, and therefore it may not be sustainable in the long term for high debt households to consume more because they expect faster house price growth.

Higher income expectations is perhaps the more convincing explanation for why high debt households consume more over a long period because differences in income growth between high and low debt households are more likely to be realised than differences in house prices. If these explanations are correct, they imply that changes in our estimate of how debt affects consumption reflect larger changes in income or house price expectations among high debt households.

Overall, it is hard to distinguish between the two explanations that either high debt levels encouraged households to spend more prior to the financial crisis and then led them to make larger revisions afterwards, or alternatively, that that households who took on high debt became increasingly optimistic about their future income prospects prior to the crisis and made larger revisions afterwards

6.2 *Evidence on why indebted households have cut spending since 2007*

The LCFS microdata is able to help us distinguish the fact that indebted households have made larger adjustments to their spending since 2007, but it cannot help to tell us why that is the case. But there is some evidence from the Bank of England/NMG Consulting survey that may be able to help. In the 2012 survey, around a third of the 4000 households interviewed said that they had cut back spending because they were concerned about their debts. Whilst this survey is more timely than our LCFS data, a net balance 12% of households reported that they have become more concerned about their debts since 2010, indicating that the results may be applicable to the post-crisis period more widely.

The NMG survey also contains questions about income shocks and about income uncertainty.

² Assuming that two households are halfway through their working life of 40 years and have identical income in the first half, one household would need to have income growth of 0.4% a year more in the second half of their working life to give them 2% higher lifetime income than the other. That would imply they consume 2% more throughout their life, which is the same effect on consumption as we estimate there is from having a debt to income ratio of 3 rather than 2. The required differences in expected future income growth are smaller the longer a household has to accumulate that extra income.

³ For example, if a household buys a house for 3 times their income and expects their house value to increase by 0.7% a year more than the interest rate on their mortgage, they expect to have 2% a year more income to consume than an otherwise identical household who bought at double their income and expects house prices to rise in line with mortgage rates. Assuming that household spend that 2% extra income, that is the same effect on consumption as we estimate there is from having a debt to income ratio of 3 rather than 2.

These questions allow us to test whether experiencing income shocks makes households more likely to become concerned about their debt and cut their spending or whether highly indebted households are more likely to have seen their income turn out lower than expected, which could be consistent with them having unrealistic expectations.

Evidence from the NMG survey is consistent with high debt levels amplifying the impact of income shocks for indebted households. Unsurprisingly, mortgagors who had cut spending in response to debt concerns tended to have higher debt to income ratios than other mortgagors (Table 1). But they were also much more likely to say that they had experience a negative shock to their income over the past year that they expected to persist and to be concerned that their income could fall over the next year. Mortgagors being much more likely to cut spending in response to debt concerns if they suffer an income shock or are uncertain about future income is consistent with our earlier econometric result that high debt household have spent less since 2007 being explained by high debt levels causing households to spend less.

Table 1: 2012 NMG survey - mortgagors' responses to debt concerns and reports of lower or uncertain future income

	Reduced spending in response to debt concerns	
	Yes	No
Median debt to income ratio	2.4	1.7
Proportion that experienced a negative income shock that they expect to persist ^(a)	32%	12%
Proportion who are think that a fall in income is quite likely over the next year ^(b)	30%	15%

(a) Questions: 'Has your household received more or less money, from both work and non-work sources, over the past twelve months than you would have expected this time last year?' and 'Are you treating this unexpected decrease in money received by your household as a decrease that is likely to persist?'

(b) Question: 'To the best of your knowledge, how likely is that your household income will fall sharply over the next year or so (for example, because you or someone in your household is made redundant)?'

There is no evidence from the NMG survey that households with high debt levels are more likely to have suffered income shocks or be uncertain about their future income (Table 2). High mortgagors are equally as likely as low debt mortgagors and non-mortgagors to have suffered an income shock over the year between 2011 and 2012 or be uncertain about their future income (Table 2). If high debt households had made larger revisions to their expected future income then we might have expected to see a greater proportion report that they had experienced an income shock.

Table 2: 2012 NMG survey - reports of lower or uncertain future income

	Mortgagors		Non-mortgagors
	Debt to income<2	Debt to income>2	
Proportion that experienced a negative income shock that they expect to persist ^(a)	21%	21%	20%
Proportion who are think that a fall in income is quite likely over the next year ^(b)	22%	21%	23%

(a) Questions: ‘Has your household received more or less money, from both work and non-work sources, over the past twelve months than you would have expected this time last year?’ and ‘Are you treating this unexpected decrease in money received by your household as a decrease that is likely to persist?’

(b) Question: ‘To the best of your knowledge, how likely is that your household income will fall sharply over the next year or so (for example, because you or someone in your household is made redundant)?’

Overall, the evidence from the NMG survey is supportive of the hypothesis that indebted household have made larger adjustments to spending since 2007, although these conclusions are tentative and do need to be treated with some caution. These results are consistent with high debt levels amplifying the impact of shocks to income. We do not find evidence that high debt households are more likely to have received lower than expected income, but this is limited by the fact that the question only asks about changes in income over the past year, ie between 2011 and 2012 and not since the start of the financial crisis. However, the 2013 survey will address this issue by asking households whether they have become worse off since 2006. This work still also does not help to tell us whether or not debt caused consumption to be higher during the period before the crisis.

7 Conclusion

We find that households with high debt levels have, on average, consumed more than those with low debt. Whilst increases in debt in the decade prior to the financial crisis were largely matched by an accumulation of financial assets, debt may have provided some support to consumption over the period when debt rose rapidly. But since 2007, the extent to which high debt households consume more has fallen back, which will have weighed on consumption growth over this period. We find evidence that weaker consumption among high debt households since 2007 may be related to debt amplifying the impact of income shocks.

Appendix:

Table A: Regression estimates

Dependent variable:	Non-housing consumption			Durable C	Non-dur. C
	(1)	(2)	(3)	(4)	(5)
LEVERAGE RATIO	0.021***	0.029***	0.023***	0.028***	0.031***
<i>Year interactions</i>					
1992		-0.022***		-0.015	-0.026***
1993		-0.032***		-0.022*	-0.035***
1994		-0.021***		-0.009	-0.028***
1995		-0.027***		-0.023*	-0.033***
1996		-0.020***		-0.013	-0.022***
1997		-0.011*		0.008	-0.016***
1998		-0.017**		-0.005	-0.020***
1999		0.001		0.022*	-0.008
2000		-0.004		-0.001	-0.007
2001		0.017**		0.022*	0.011*
2002		0.012*		0.025**	0.002
2003		0.014**		0.034***	0.002
2004		0.004		0.015	-0.004
2005		0.002		0.009	-0.002
2006		-0.003		-0.001	-0.004
2008		-0.012**		-0.010	-0.013**
2009		-0.023***		-0.027**	-0.022***
2010		-0.009		-0.011	-0.011*
2011		-0.018***		-0.016*	-0.019***
POST 2007 DUMMY			-0.010***		
CONTROLS	YES	YES	YES	YES	YES
Observations	103,309	103,309	103,309	102,972	103,309
R-squared	0.587	0.588	0.587	0.450	0.548

*** p<0.01, ** p<0.05, * p<0.1. Dependent variable is log of real household consumption, net of housing expenditure. Leverage ratio is the total amount of outstanding secured debt divided by household disposable income (2007 is omitted). Income is the log of real household disposable income. Controls include year & quarter dummies, income Government Office Region, age up to the 5th power, number of children and adults in the household by age group, household size, ethnicity, marital status, gender, occupation, school leaving age, tenure, date of birth cohorts & regional house prices.

References

To be added