



BANK OF ENGLAND

Speech

From retailers' paradise to shoppers' strike: what lies behind the weakness in consumption?

Dean's lecture given by

Martin Weale, External Member of the Monetary Policy Committee, Bank of England

At Cass Business School, London

Wednesday 29 February 2012

I would like to thank Matthew Corder, Robert Gilhooly and Tomasz Wieladek for research assistance, and to Justin Van de Ven for his advice on use of the National Institute Tax and Benefit model. I am also grateful for helpful comments from other colleagues. The views expressed are my own and do not necessarily reflect those of the Bank of England or other members of the Monetary Policy Committee.

All speeches are available online at www.bankofengland.co.uk/publications/speeches

Introduction


Thank you very much for the kind invitation to give this Dean's Lecture at the Sir John Cass Business School. Sir John Cass was of course a great philanthropist. His will was contested in the Court of Chancery but the dispute seems to have fared better than the later case of Jarndyce versus Jarndyce described by Charles Dickens because the modern foundation which bears his name supports a number of educational establishments including this School. The establishment of a foundation inevitably reflects the importance its founder attaches to future rather than current needs; neglect of future needs and future risks is a recipe for economic disaster. But today I would like to look at the other side of that coin – consumption. Although the most recent GDP and retail sales figures point to some improvement, one of the features of the recent period of recession and stagnation has been weakness in consumption – quite a change from the years before the crisis.

Household consumption (at a little over 60% of GDP) is much the largest component of total expenditure and, as a result, even though it is less volatile than investment, fluctuations in consumption have an important effect on overall GDP. Economic theory suggests that consumption is strongly influenced by people's expectations of future incomes as well as being affected by current income. This means that fluctuations in consumption can be regarded as indicative of expected future movements in income; weak consumption may be an indicator of persistent weakness in income. It is therefore important that the MPC understands the drivers of consumption if it is to assess the outlook for demand and inflation.

In this speech I will outline the scale of the fall in consumption since the start of the crisis and identify some of the factors lying behind it. The fall in people's incomes since the early 2008 explains much of the decline in consumption. But I will also outline how analysis that takes account of the different shocks faced by different people or the way that they are affected differently by the same shocks (such as the adverse labour market conditions faced by young adults since the crisis, the effect of changing state pension ages, and the increased risk of unemployment) may also contribute to an understanding of aggregate consumption. I will conclude with some thoughts on how this feeds in to my overall thinking on monetary policy.

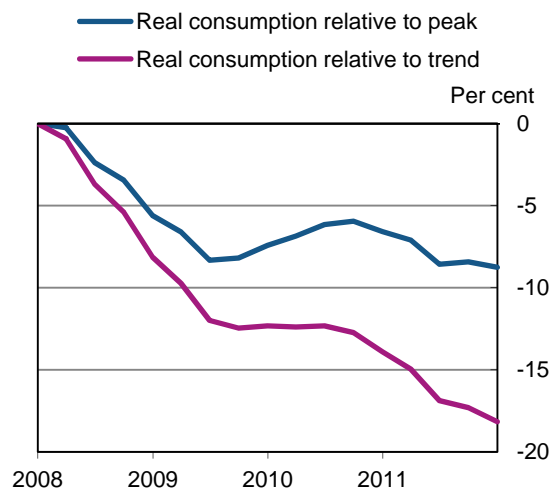
If consumption is an indicator of future income, then recent data may point to further weakness in income to come. At the end of last year, real disposable income per person had fallen by 5 per cent from its peak in late 2007, and by around 15 per cent relative to its long-run linear trend (Chart 1). But consumption per head has fallen even more than income. Real consumption per capita in the second half of 2011 was over 8 per cent below its 2007 peak, and over 18 per cent below its pre-crisis trend (Chart 2).

Chart 1: Real income per capita since 2008



Sources: OBR, ONS and Bank calculations.

Chart 2: Real consumption per capita since 2008



Sources: ONS and Bank calculations.

In very broad terms there is considerable consensus among economists about influences on consumption. There is quite good evidence that the majority of income accrues to people who plan their spending with a view to the future as well as to the present – that they seek to smooth out the effects of fluctuations in income. Modigliani and Brumberg (1954) developed the life-cycle hypothesis, which, in essence, assumes that people plan their consumption over their lives in much the same way as they allocate their consumption between different types of good at any point in time. This leads to the conclusion that short-term fluctuations in income have little effect on consumption while those changes which are expected to be permanent have a powerful effect.

But within this general framework there is plenty of debate. What happens if risk increases? How is consumption behaviour affected by the support offered by the state benefit and pension system? What might the role of fluctuations in credit availability be in influencing consumption? While one might hope to have a sophisticated framework capable of explaining fully the effects of the major influences on consumption I do not think applied economics has yet reached the point where that is possible. In particular, notwithstanding recent progress (e.g. Waldron and Zampoli, 2010, Attanasio, Leicester and Wakefield, 2011), there is as yet no fully-coherent and convincing analysis of the interaction between house purchase, mortgage credit and consumption. Nevertheless, what I can do in this talk is to air a number of factors relevant to consumption analysis and quantify the influence of some of these.

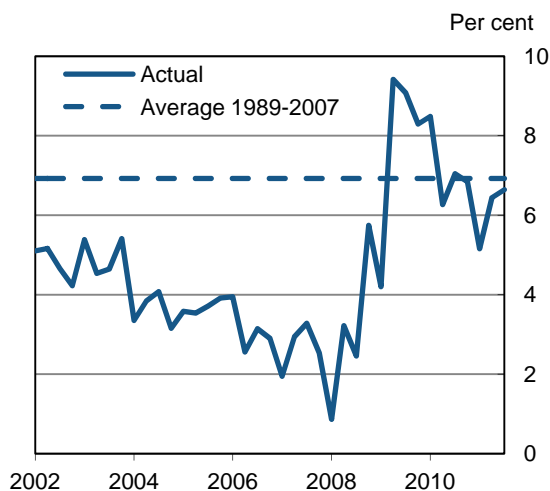
Consumer Spending and Monetary Policy

Monetary policy is effective, in part because of its impact on consumer demand. As you know, the Bank Rate was reduced from a level of 5 per cent in the summer of 2008 to the level of ½ per cent today. Rates

faced by households have not fallen to the same extent, but average rates paid on deposits have dropped by around two percentage points and those charged on mortgages have fallen by 2½ percentage points. We have, of course, also undertaken two asset purchase programmes. The first started in 2009 and finished in January 2010. The second started in October of last year and was extended earlier this month.

As I noted in my introduction, there is broad consensus that the majority of consumers plan their spending with regard to future as well as current income. Provided people can form a sensible expectation of the resources available to them over their life-times, they can then choose the consumption plan which gives the highest possible level of welfare over their remaining expected life-time while at the same time deciding on any legacies they would like to leave. This leads to the conclusion that they save when income is temporarily high and borrow or reduce their savings when income is temporarily low. In this framework consumption growth depends on the real interest rate. Low interest rates reduce the incentives to save and encourage consumers to consume more now at the cost of consuming less in the future. So current consumption is raised but future growth rates of consumption are reduced. This is one of the important channels by which monetary policy affects economic activity and thus inflation.¹

Chart 3: Household savings ratio



Source: ONS

The Office for Budget Responsibility (OBR) expects real income per capita to be almost 20 per cent below pre-crisis trend by 2017. After a large adverse shock to expected future income, consumption should fall to ensure, over time, income can pay for expenditure. As Chart 2 shows, consumption fell in 2008. But, if consumers had foreseen the full impact of the financial crisis and rising commodity prices on their future income, then, without any policy change or expectation of a rapid return to normality, consumption would initially have fallen even more sharply than it actually did, and the savings rate would have temporarily risen even more abruptly. This extra saving would of course, have added to wealth, so that the eventual fall in disposable income would have been lower than that indicated. But the wealth offset would have been relatively small. If the full effect of the crisis on income had been foreseen at the time, then this simple analysis would have suggested that consumption should have fallen by a proportion not much smaller than the eventual fall in income.

In fact the adjustment has been relatively gradual as Chart 2 shows. There are a number of possible explanations. First of all, it has only become clear gradually just how large the adverse shock to incomes

¹ For more on the transmission mechanism of monetary policy see: <http://www.bankofengland.co.uk/publications/other/monetary/montrans.pdf>.

has been since the financial crisis. Secondly, people may adjust their spending more slowly than the framework I described suggests, because they like to follow their habits - an issue originally raised by Stone and Rowe (1958) - or perhaps simply because they try to keep up appearances (Abel, 1990). But thirdly, the Monetary Policy Committee has acted to support output so that inflation did not undershoot the Bank's target of 2 per cent. By rapidly cutting Bank Rate, the MPC encouraged consumers to replace future consumption by spending now.

A simple back of the envelope calculation can show what might have happened, had consumers anticipated the path actually followed by household real interest rates since the crisis and were those rates in future to move in line with market expectations up to 2015 Q1 but return to pre-crisis levels at the end of 2018.² Then consumption would have been supported initially by at least four percentage points compared with what might have happened in the absence of a policy relaxation. Indeed, consumption is now at a position broadly consistent with market expectations of the interest rate and the income projections I described above. At least this is the case if I expect the savings ratio to settle at the average of 1989 to 2007 which is just under 7 per cent. But these calculations are based on the assumption that the influence of the interest rate on consumption is that identified by van de Ven (2011);³ this effect is lower than some other estimates have suggested. So perhaps consumption is on the low side of what might be expected. Of course this is also true if consumption is based on expectations of future income more favourable than those projected by the OBR.

On top of this, of course our policy of asset purchases has provided support to consumption through its effects on asset values, and in particular by providing some support for prices of financial assets and property. The analysis I have described gives less role to spending out of wealth than some empirical evidence suggests. (See e.g. Barrell and Davis, 2007, Campbell and Cocco, 2007, but for a contrary view see Attanasio, Blow, Hamilton and Leicester, 2009. Lettau and Ludvigson, 2004, argue that there is little spending out of transitory fluctuations in wealth but substantial spending out of so-called permanent wealth.) But I would like to explore some of the factors which may be weighing down on consumption before returning briefly to the issue of housing and wealth.

Age Effects

The framework that I have described assumes that all consumers are alike. In practice, the differences between them might have important consequences at the aggregate level. So the first issue I would like to investigate is how far people of different ages have had different experiences recently. To study consumption at the disaggregated level it is necessary to rely on the data contained in the Living Costs and

² The market yield curve suggests that real rates are expected to remain appreciably lower than before the crisis almost indefinitely. But the same effect (Deaton, 1992) can be achieved with the assumption that there has been an increase in the perceived uncertainty of consumption, so that, risk adjusted, normal rates of consumption growth are resumed then. Without such a risk premium the relationship between consumption levels and the market yield curve is very difficult to explain.

³ The assumed inter-temporal elasticity of substitution is 0.34.

Food Survey (LCF) and its predecessors (the Expenditure and Food Survey and the Family Expenditure Survey), while the Labour Force Survey tells us what has been happening to people's wages.

In Chart 4 I show the growth in the real wages of people in full-time paid employment. We can see a very clear picture from this. Before the crisis, the growth in wages of young adults was slower than that of old workers. But since the crisis this differential has become much more marked with real wages of young adults falling markedly more than those of people close to retirement. To the extent that these movements lead to similar changes in expected life-time wage income, this amplifies the impact of weak incomes on consumption compared with a situation where the weakness in wages affected people of different ages similarly.

At the same time, of course, we know that the unemployment rate of people aged 18-24 has risen from 12 per cent in early 2008 to around 20 per cent nowadays. For those aged twenty-five to thirty-four the unemployment rate has risen from 4½ per cent to around 8.3 per cent while for the whole population the increase has been from 5.2 per cent to 8.4 per cent.


It is possible, of course, that these two observations are inter-connected. Wages of young adults may be particularly depressed because there is a disproportionate number of them looking for work. It might be hoped that the high unemployment rate has little impact on future earning potential and also that, as labour market conditions improve, so the wages of young adults pick up. But research by Gregg and Tominey (2005) and Arulampalam (2001) suggests that unemployment when young may have a persistent effect on wages ten and twenty year later, particularly if the duration of unemployment is long (Gregory and Jukes, 2001). This suggests that the overall wage outlook for young people may be even worse than would be suggested if the effects shown in Chart 4 simply persisted. It is certainly the case that the past few years have not been kind to young adults.

If I look at consumption as measured by the LCF and allocate household consumption across individuals I also find that the consumption of young people has been weak relative to that of old people (Chart 5). One might also speculate that the rapid growth in consumption of people over fifty between 2003 and 2007 was influenced by the house price boom at that time (Attanasio, Leicester and Wakefield, 2011). And high asset prices may have continued to support the consumption of people aged sixty-five and over since 2007.

Why might issues of this type matter even if we are concerned only with aggregate consumption? Young people have not had much opportunity to build up wealth, so their consumption will depend mainly on expectations of future labour income. Old people, close to or after retirement will, by contrast, expect to finance most of their subsequent consumption out of accrued wealth (including pension rights and entitlements to state pensions and other benefits) and will not find themselves nearly so much affected by adverse movements in either wage rates or employment opportunities. Indeed simulation of a simple life-cycle model suggests that this differential effect could augment the impact of the squeeze in real wages

on consumption by about one percentage point as compared to a situation where the same average fall in real wages had been spread uniformly across people of all ages.

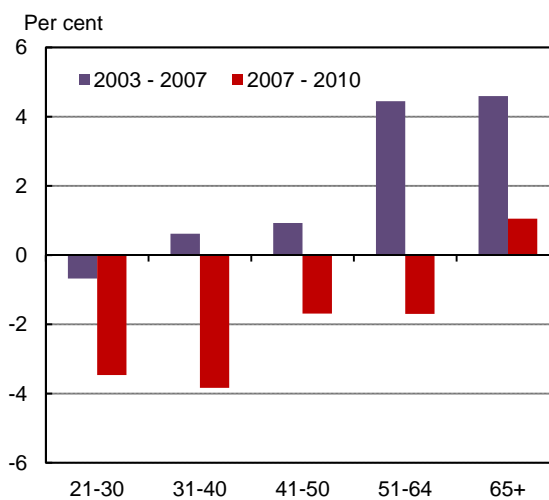
Chart 4: Average annual real gross wages growth rates^(a) by age group



Sources: LFS and Bank calculations.

(a) Individuals in full time employment.

Chart 5: Average annual real consumption growth rates by age group



Sources: LCF, EFS, FES and Bank calculations.

A Structural Analysis of Disturbances to Consumption

The standard modelling tool used in central banks relies on the assumption that there are representative consumers. In its simplest form it is assumed that there is a single individual whose spending decisions represent the choices of all the families in the economy. These models can be extended to address the possibility that a constant proportion of income accrues to households who spend their income as it comes in. But, beyond this there is a range of factors which they cannot address.

One influence on saving missing from standard models is uncertainty. Economists from Keynes (1930) onwards have drawn attention to the precautionary motive for saving; people save to provide some protection from the vagaries of the future; if there has been any increase in the perception of uncertainty, then saving is likely to increase. As I noted in footnote 2, it has much the same effect as an increase in the interest rate. Could an increase in uncertainty be a factor depressing consumption relative to income and if so might it have further to run?

A second point is that a model based round a representative consumer cannot represent what I suggested earlier might be the main long-term motive behind saving, saving for retirement. We know that there have been a number of changes to the social security system announced over the last few years, of which increases in the state pension age have been perhaps the most prominent. That is likely to affect saving behaviour, and also labour supply, but a representative agent model can hardly be expected to show this.

A third influence I would like to say something about is the effect of the availability of credit. On the face of it a growing economy in which people cannot borrow is likely to have more saving than one in which borrowing is possible. Consumption has to wait until people have the money to pay for it and that means it comes later than if credit is available. But the effect of credit availability is not totally disconnected from the impact of uncertainty. In a world in which credit is available, people whose incomes are unexpectedly, but, they believe, temporarily, low, can address the problem by borrowing. In the absence of credit they can smooth out the effects of such disturbances only by drawing on past saving. So the magnitude of the precautionary motive is also likely to be affected by the readiness and the terms on which people expect credit to be available.

Deaton (1992) was the first to succeed in developing a useable model of the behaviour of consumers whose consumption differs because of their circumstances, but not their preferences. Since then a wide variety of authors has developed his approach and, as computers have speeded up, so the richness of the models has increased. The underlying assumption behind them is that people plan their consumption and their labour supply so as to maximise the level of welfare they can expect to enjoy for the rest of their lives. These decisions are made taking account of the fact that the future is uncertain; incomes can be forecast only imprecisely and people cannot predict the date of their death.

Such a model has been developed by my colleagues at the National Institute (van de Ven, 2011) primarily with the purpose of looking at the impacts of taxes and benefits on people's behaviour (Sefton, Weale and van de Ven, 2008, Sefton and van de Ven 2009). At present there are many obstacles to the use of such models as part of a suite of tools for regular forecasting and analysis. The results are obviously dependent on the parameters of the model, a feature which cannot be avoided, and an exploration of the sensitivity of the findings to key parameters would itself be a very substantial exercise. But they do allow us to look at some of the macro-economic issues I raised above, at least to some extent.

Risk, Unemployment and Saving

The future is self-evidently uncertain. A large risk that people face is that they will find themselves unemployed and the magnitude of this risk is likely to affect people's saving behaviour. Simulation of household behaviour can take account of this, by imposing an exogenous probability that people simply cannot find work. An economy in which this probability is high can then be compared with one in which it is low, and the difference between the two can be used to estimate the effect of an unexpected permanent increase in unemployment risk on saving behaviour.

In broad terms an increased risk of unemployment should be expected to have two main effects. First of all average incomes will be lower, simply because fewer people on average will be working. This will reduce saving in absolute terms. Saving as a proportion of income may also appear reduced because, in the partial

analysis I present, state retirement benefits are assumed to remain unchanged. Thus the need for pension saving declines more than in proportion to the decline in labour income.

Offsetting this, in an economy with a higher probability of not being able to find a job, people will want higher savings, so as to have an increased buffer should they be unemployed. Young people will be disproportionately affected in the way. People near the end of their working lives may come to the conclusion that the buffer they have built up can now be run down a bit. So in a steady state economy with high unemployment, young people may save more than they otherwise would, while people late in working life may save less than they otherwise would.

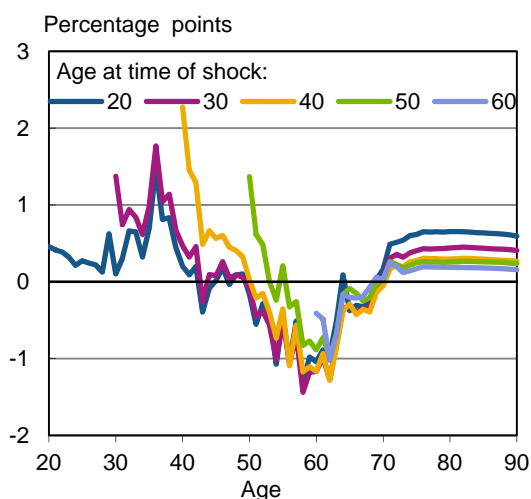
These effects can be seen if we compare a base economy with one in which the risk of not being able to find a job is increased by three percentage points. Total employment income is reduced to nearly the same extent although the buffering arising from the tax and benefit system means that disposable income is reduced by less. Non-pension wealth is increased but total wealth declines.

But a comparison of two different economies does not tell us about the short-term impact of an increase in unemployment which is believed to be permanent. People who have saved in the expectation of a low unemployment risk and then find themselves facing a high unemployment risk which they expect to persist for the rest of their working lives will react differently from those who have always lived with a high unemployment risk. And the impact is bound to depend on whether they are close to retirement or have just started working.

In Chart 6 I show the aftermath of a three percentage point unemployment shock on saving at different ages. The apparently erratic movements are a consequence of the solution method, which always involves a compromise between precision and computing time. But the lines do allow us to tell a clear story. Young people would like to consume more than their income and therefore do not raise their saving in response to an increase in unemployment risk by very much. They anticipate that they will be less constrained by the time they reach their mid-thirties and trade off some of the desired protection in the buffer offered by extra saving until then. This is also true, to some extent, if thirty-year olds are affected by the shock. But forty and fifty-year olds react sharply to the extra risk that they face; the impact on their saving is at a maximum in the immediate aftermath of the shock.

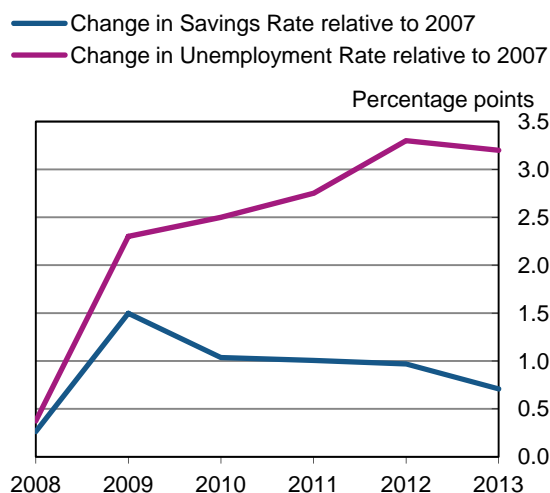
Late in working life, we can, as predicted, see that people feel they now have less need of precautionary balances, and saving falls below where it would have been in the absence of the shock to unemployment. Beyond about seventy, however, saving seems to increase again. This reflects the fact that, in the aftermath of the shock, people plan for lower retirement consumption than they would have chosen in its absence. As a consequence the decumulation of wealth is slowed and this naturally enough appears as an increase in net saving.

Chart 6: Effects on Savings Profiles of a Permanent Increase in Unemployment^(a)



Source: National Institute Benefit & Tax Model .
 (a) The graph shows the change in saving measured as a proportion of disposable income in the absence of a shock to unemployment.

Chart 7: Effects of Unemployment on Savings^(a)



Source: National Institute Benefit & Tax Model and OBR.
 (a) The graph shows the change in saving measured as a proportion of disposable income in the absence of a shock to unemployment.

These effects can be put together to give an overall indication of the short-term impact of a permanent rise in unemployment on the aggregate savings rate, by weighting the above figures by the disposable income of households headed by people at each age. If I make the assumption that each of the changes to unemployment since 2007 have been unanticipated and is expected to be permanent, project forward the unemployment rate using the projections made by the OBR in its November forecast, and then cumulate the effects shown in Chart 6, the effect of this on saving, relative to its pre-crisis position is shown in Chart 7.

We can see, first of all that the effects of the rise in unemployment can be substantial, and secondly that, although the influence of a rise in unemployment fades with time. Nevertheless, these effects are less marked than those identified by Modey, Ohnsorge and Sandri (2012). The reason for this is almost certainly that they assume the sole risk people face to labour income is that of unemployment. In other words they neglect the normal labour market churn which is a substantial source of uncertainty to many people and thus overstate the proportionate impact of increased unemployment.

More importantly, if unemployment falls off next year as forecast by the OBR, the upward pressure on the savings rate from this source is likely to ease. Of course, if people's expectations adjust slowly or if consumption is based partly on habit, then these effects may be smoothed out so that, in particular, the drop in the savings rate next year may be less marked.

However we should avoid drawing the conclusion that movements in unemployment necessarily exert a strong pro-cyclical influence on demand through this route. The impact is shown on the assumption that

people expect the increase in unemployment to be permanent. Such expectations might not be rational in a typical economic cycle.

A Rising State Pension Age and Social Security Reform

The state benefit system is likely to have a substantial impact on household savings behaviour. There are two reasons for this. First of all, the payment of state pensions to old people reduces the need to save for retirement. The retirement pension is not large compared with, for example, the median income of someone in full-time work. But, nevertheless, it is a close to universal benefit and provides a substantial resource to people in old age that they would otherwise have to do without, or provide by means of saving while at work. Secondly, the system of benefits available to people while of working age does provide some sort of back-up in the event of unemployment and other unfortunate developments such as ill health. In its absence the level of precautionary saving would be higher.

The social security system as a whole is undergoing substantial change with the introduction of the Universal Credit for people in work. At the same time increases in the age at which state pensions become available have been announced in two phases. The 2007 Pensions Act set an eventual pension age of sixty-eight with earlier rises to sixty-six and sixty-seven. The 2011 Pensions Act accelerated these increases. The pension age is set to rise to sixty-six over a two-year period centred approximately on 2019 and to sixty-seven over a two-year period centred on 2027. The date of the increase to sixty-eight remains at 2045 as specified by the 2007 Act.

The full range of changes taking place to the social security system means that it is, to say the least, not straightforward to analyse their impact on savings behaviour. I would like to limit myself to exploring the effects, on their own, of the increases in the state pension age. To simplify the analysis I am exploring the short-term effect of an unanticipated announcement of an increase in the state pension age from sixty-five to the pattern of pension ages which emerged after the changes in 2011.

In Chart 8 I show the simulated effect of this on the saving behaviour of people of different ages chosen to reflect the dates of changes to the state pension age. The chart shows the change in saving as a percentage of mean disposable income by age before the change. It indicates the effect on people from their age at the time of the announcement to age sixty-five. A striking feature of the simulations is that the announced increase in the pension age has little impact on saving until people reach their early-40s. The explanation for this is two-fold. Young people on average would like to consume more than their incomes because they, at least on the basis of past experience, can expect their incomes to rise with age. People in their thirties and early forties are likely to save. But at this age it is likely that saving behaviour is dominated by the precautionary motive. People wish to build up a buffer of savings so as to have something to fall back on if, for example, they become unemployed.

This observation is entirely consistent with the analysis presented by Carroll (1997). He suggested that saving for retirement does not become important until people reach their mid to late forties.⁴ Chart 8 points to a slightly earlier start, but later in life we can see that the deferment of the retirement age has a substantial impact on saving behaviour. This effect is more marked for those currently in their forties than for those currently in their fifties mainly because the former face a pension age of sixty-seven while the latter will be able to draw their pensions at sixty-six.


One reason why the increased saving takes place relatively late in working life is that part of the mechanism which makes increased saving possible is postponed retirement; the simulations reflect the consequences of increased labour supply from people in their late fifties and early sixties.

We can form an overall estimate of the effect of the change to retirement dates by applying the impact effects shown in Chart 8 to estimates of disposable income for households whose heads have ages close to the ages shown. Thus it is assumed that the fifty-five year old household represents the behaviour of households with heads aged fifty-two to fifty-seven, whose retirement age is sixty-six. The behaviour of younger households is represented by the impact effects shown for the younger ages as appropriate.

These calculations suggest that the effect of announcing a change from retirement at age sixty-five to the pattern defined by the 2011 Act would be to raise the savings rate by about 1.1 percentage points. Eventually, of course, as the people whose saving is increased actually reach sixty-five they will start to draw down on their savings as a partial replacement for the pensions they no longer receive. So the eventual impact on saving will be small. But, looking at the impacts after one, two and three years, the simulations do not suggest that the effects are likely to decline, because, as the people affected move closer to retirement, their savings rates tend to rise rather than to fall. For quite a while we are likely to see savings rates around one percentage point higher than they would have been without the change.

I am not, of course, claiming that the overall effect of all the changes to the pension and benefit system is to raise the savings rate in this way. Other important changes include the introduction of the Universal Credit

Chart 8: The Effect of Increasing the State Pension Age on Saving^(a)



Source: National Institute Benefit & Tax Model.
(a) The graph shows the change in saving measured as a proportion of disposable income in the absence of a shock to unemployment.

⁴ It should be noted that this result does not depend on the combination of a discount rate of 10 per cent per annum and zero real interest rate assumed by Carroll. The model used here has a discount rate of just over 3 per cent per annum and returns of 2.7 per cent per annum or 4 per cent per annum depending on whether savings are liquid or not.

and the commitment to increasing the retirement pension annually by highest of wage growth, price increases or 2½%. But the exercise does demonstrate that, with plausible assumptions about people's behaviour, changes to social security can have a non-trivial influence on saving and might lead us to expect a savings rate higher than we have seen in the past.

Housing, Wealth and the Availability of Credit

My initial analysis was conducted with respect to disposable income. Since wealth is the capitalised value of property income and since I assumed that real interest rates will return to their pre-crisis levels, there was an implicit assumption that asset holdings *per capita* eventually fall in line with disposable income. The reason for looking at the problem in this way rather than through the more conventional framework which focuses on non-property income and initial wealth holdings is that returns on major asset holdings, housing and equities, have not fallen in line with the short-term interest rate, even after adjusting for credit spread effects. And one has to remember that, before the crisis, most indicators suggested that house prices were unusually high. It is perfectly possible that the decline in real house prices since the crisis has been an important factor in bringing the savings ratio much closer to its historic average. At the same time, it is noteworthy that, while real house prices have fallen by over a quarter since their peak, they are, relative to income, still 8 per cent above their long-run average. This could largely be explained by low real short-term interest rates. Or it could be explained by other offsetting factors. For example, tight credit could be depressing house prices while they may also be receiving substantial support as a consequence of the effects of the Bank's asset purchase programmes on long-term interest rates.

Some studies suggest that housing wealth exerts a stronger influence on spending than do other forms of wealth; the impact depends on a range of factors including the extent to which people feel they should leave their houses to their descendents. But if instead of looking just at house prices I look at overall household sector wealth the picture is quite striking. Averaged over the period 1989-2007 households held net wealth equal to 5.72 times their income. The most recent estimates suggest that that ratio is currently 6.64. This suggests that to the extent that people spend more than just the property income earned by their wealth, we should expect to see consumption supported rather than squeezed by balance sheet effects. But it is hard to see this effect⁵ amounting to more than one to two percentage points since, in working with reference to disposable income, I am already taking account of spending out of the income generated by that wealth.

Might there nevertheless be separate influences on consumption from movements in the availability of credit? Access to unsecured credit means that young people, who might reasonably expect their incomes to rise fairly sharply early in working life, can borrow to smooth out their consumption over their lifetimes. Reduced access to credit limits their consumption when young and raises it when they age, but it does not,

⁵ If I assume that the real return on wealth after tax is 3 per cent and that the propensity to spend out of wealth is 4 to 5 per cent. The assumption about the savings rate means that nearly all the income from wealth is already assumed spent. Such a propensity to spend out of wealth is more appropriate for permanent than transitory wealth (Lettau and Ludvigson, 2004).

of course, lead to a corresponding increase in consumption of today's old people. Access to unsecured credit or the ability to use a home as collateral reduces the need for precautionary saving (Lustig and van Nieuwerburgh, 2005). For both these reasons, both a tightening of the conditions on which credit is available and a fear, whether well-founded or not, that people will not be able to obtain credit in the event of an adverse shock to their income, will lead to an increase in saving.

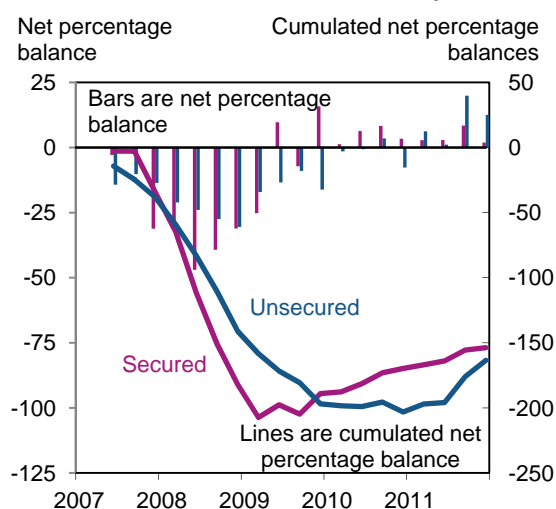
The framework I used to explore unemployment risk and the effects of delayed retirement is of only limited use to explore credit issues; it does not fully explain the observed level of unsecured borrowing.

Chatterjee, Corbae, Nakajima and Rios-Rull (2007) in an impressive piece of work, developed a model in which the threat of penalties associated with bankruptcy generated an equilibrium in which most people repaid their unsecured debts while the risk premium on unsecured lending was endogenous. But there are questions how well their model represents the various risks people face and I have not seen it used to attempt to account for the factors which might lead to a reduction in unsecured indebtedness of the type which has actually occurred.

The problems which arise in the use of this sort of framework to examine the effects of changed perceptions about access to mortgage credit are, if anything, even more acute and simplification is needed to make any progress. Thus Attanasio *et al.* (2011) assume that home owners can choose between only two types of homes, flats and houses. All houses have the same price and flat prices are a fixed fraction of house prices. These simplifying assumptions mean that they can represent the point that people make housing transactions only infrequently and that the dominant effect of high house prices on consumption comes through its influence on the consumption of old people. As I noted, this is a possible factor behind the pattern shown in Chart 4.

The issue of credit availability, however, cannot be ignored. The Bank of England's Credit Conditions survey suggested that the availability of both secured and unsecured credit to households fell sharply in 2008 (Chart 9), with the decline in the availability of unsecured credit stretching in to early 2010. Over this period the ratio of unsecured debt to income has fallen steadily from an average of 26.8 per cent in 2007 to 23.9 per cent by late 2011. The ratio of secured debt to income has fallen from an average of 120 per cent of income in 2007 to 114 per cent of income by late 2011. On the other hand the decline in house prices means that, relative to housing wealth, secured debt has risen by around 6 per cent.

Chart 9: Household credit availability^{(a)(b)}



Source: Bank of England
(a) Net percentage balances are calculated by weighting together the responses of those lenders who answered the question. The bars show the responses over the previous three months.
(b) A positive balance indicates that more secured credit is available.

The survey suggests that the tightness of credit has not reversed since the early days of the financial crisis. This may affect consumption in several ways. First, one would expect the impact of unexpected difficulty in access to credit to have a sharp impact on spending in the short term and a fading impact in the longer term. Secondly, households' concern that they may not have access to credit, when they might find it useful, is likely to make aggregate consumption more sensitive to income and less sensitive to interest rates than would otherwise be the case. As I have explained above, there is no straightforward way of quantifying this effect satisfactorily.⁶ Nevertheless, any easing of credit conditions would be likely to support consumption.

Where do we go from here?

The main driver of consumption I have identified is income. This is hardly an original point, but I come across a remarkable number of people who stress what they call consumer confidence and lose sight of the fact that spending has to be paid for somehow. The large fall in national income following the crisis therefore accounts for the bulk of the fall in consumption we have seen over recent years. I have also identified a number of additional influences which are likely to depress consumption relative to pre-crisis experience: the hit to income that is skewed towards the young; changes in state pension age; a greater risk of unemployment and tighter credit conditions. On the other side, households' high level of wealth relative to income could have helped support consumption.

Turning to the most recent data, consumption appears to be growing again. The ONS estimate that private consumption grew 0.4% in Q4 – the first increase in consumption since 2010 Q3 – though overall GDP fell 0.2% as a result of weak business investment and reduced stockbuilding. Indicators of the state of the economy this year have been more positive for overall GDP growth, and the retail sales figures suggest consumption may be continuing to rise in Q1. At the same time we need to remind ourselves that the signals they provide are imprecise and that a few swallows do not make a summer. The distortions resulting from the holiday for the Queen's Diamond Jubilee mean that the data, as they emerge, will be harder to interpret than is usually the case.

Looking ahead, the forecast in our *Inflation Report* suggests that reasonably normal economic growth is not likely to resume until late this year. If evidence builds that the OBR's projections for future incomes are too pessimistic, then we might expect to see more rapid growth of consumption and the shoppers' strike will lose some of its force. And some of the factors weighing on consumption are likely to fade as unemployment eventually starts to fall back or as the availability of credit increases again. But other factors that have been depressing consumption are likely to persist, so there is a risk that the savings rate will settle at a level higher than the average of past experience. Since the actual savings rate is currently just below the average of past experience (Chart 3), this points to limited room for a rebound in consumption growth. And while low expected interest rates can support the level of consumption, because they bring consumption forward, they

⁶ The profile I have used for future household interest rates already assumes that the excess margin faced by households over Bank Rate declines with time. But availability of credit is a separate matter.

could, for the same reason, lead to a lower growth rate of consumption over the medium term. After all, people cannot spend their money twice. So on balance, what I have described represents a downside risk to the medium-term growth projections in our recent *Inflation Report*.

Such a downside risk may not, of course, materialise and does not automatically translate into a presumption that inflation will be markedly lower than the *Report* indicates for two reasons. First, even if consumption does grow less rapidly than income, it is possible, that improving export demand or rising business confidence and reviving investment will nevertheless support the economy; indeed the most recent CBI survey suggests a much improved export picture. Secondly, weaker overall demand does not necessarily translate into below-target inflation in the medium term.

Indeed, although the economy has recently been weak, inflation remains above target. The inflation rate has fallen sharply as the effects of the VAT increase and some of the fuel price increases have dropped out of the numbers. But, on a seasonally-adjusted basis, the monthly rate of inflation, which fell below a level compatible with the target at the end of last year, seems likely to be higher than is consistent with the target in the first half of this year. The price of oil is a particular worry. The further out we go, of course, the more uncertain things become. Nevertheless, this does suggest a risk that there may be more persistence to inflation than one might expect at a time of rising unemployment and weak demand. Further ahead, there remains a risk that an eventual return to more normal economic conditions will be associated with increased wage pressures. Judgements about the magnitude of this are inevitably uncertain and it is hard to avoid the sense of an additional upside risk.

I argued, after our November forecast, that one could be less pessimistic about inflation than our forecast showed – pessimistic in the sense of seeing it less far below target at the two-year horizon – but still see a case for further quantitative easing once the programme we announced in October came to an end. At present for the reasons I have outlined above, I do not think there is likely to be a further case once our current programme is complete. Indeed the yield curve suggests that an increase in Bank Rate is not fully priced in until mid-2014. But, obviously, if the very real risks I see about inflation do materialise, then it is perfectly possible that the first rise will come earlier than that.

References

- Abel, A B (1990)**, "Asset Prices under Habit Formation and Catching Up with the Joneses", *American Economic Review*, Vol. 80, pages 38-42.
- Arulampalam, W (2001)**, "Is Unemployment Really Scarring? Effects of Unemployment Experiences on Wages", *Economic Journal*, Vol. 111, No. 475, pages F585-F606.
- Attanasio, O P, Blow, L, Hamilton, R and Leicester, A (2009)**, "Booms and Busts: Consumption, House Prices and Expectations", *Economica*, Vol. 76, No. 301, pages 20-50.
- Attanasio, O, Leicester, A and Wakefield, M (2011)**, "Do House Prices Drive Consumption Growth? The Coincident Cycles of House Prices and Consumption in the UK", *Journal of the European Economic Association*, Vol. 9, No. 3, pages 399-435.
- Barrell, R and Davis, E P (2007)**, "Financial Liberalisation, Consumption and Wealth Effects in Seven OECD Countries", *Scottish Journal of Political Economy*, Vol. 54, No. 2, pages 254-267.
- Campbell, J Y and Cocco, J F (2007)**, "How Do House Prices Affect Consumption? Evidence from Micro Data", *Journal of Monetary Economics*, Vol. 54, No 3, pages 591-621.
- Carroll, C (1997)**, "Buffer-stock saving and the Life-cycle/Permanent Income Hypothesis", *Quarterly Journal of Economics*, Vol. 112, pages 1-55.
- Chatterjee, S, D Corbae, M, Nakajima and Rios-Rull, J-V (2007)**, "A Quantitative Theory of Unsecured Consumer Credit with Risk of Default", *Econometrica*, Vol. 75, pages 1525-1589.
- Deaton, A (1992)**, *Understanding Consumption*, Clarendon Press Oxford.
- Gregg, P and Tominey, E (2004)**, "The Wage Scar from Youth Unemployment", CMPO Working Paper 04/097.
- Gregory, M and Jukes, R (2001)**, "Unemployment and Subsequent Earnings: Estimating Scarring Among British Men 1984-94", *Economic Journal*, Vol. 111, No. 475, pages F607-F625.
- Keynes, J M (1930)**, *Treatise on Money*, Macmillan, London.
- Lettau, M and Ludvigson, S C (2004)**, "Understanding Trend and Cycle in Asset Values: re-evaluating the Wealth Effect on Consumption", *American Economic Review*, Vol. 94, pages 276-299.
- Lustig, H N and van Nieuwerburgh, S G (2005)**, "Housing Collateral, Consumption Insurance and Risk Premia: an Empirical Perspective", *Journal of Finance*, Vol. 60, pages 1167-1219.
- Modigliani, F and Brumberg, R (1954)**, "Utility Analysis and the Consumption Function: an Interpretation of Cross-section Data" In K Kurihara (ed.) *Post Keynesian Economics*, Rutgers University Press, New Brunswick.
- Mody, A, Ohnsorge, F and Sandri, D (2012)**, "Precautionary Savings in the Great Recession", IMF Working Paper, WP/12/42, <http://www.imf.org/external/pubs/ft/wp/2012/wp1242.pdf>.
- Sefton, J, van de Ven, J and Weale, M R (2008)**, "Means Testing Retirement Benefits: Fostering Equity or Discouraging Savings", *Economic Journal*, Vol. 118, pages 556-590.
- Sefton, J and van de Ven, J (2009)**, "Optimal Design of Means Tested Retirement Benefits", *Economic Journal*, Vol. 119, pages. F461-81.

Stone, J R N and Rowe, D A (1958), "Dynamic Demand Functions: some Econometric Results", *Economic Journal*, Vol. 68, pages 256-270.

Van de Ven, J (2011), "A Structural Dynamic Microsimulation Model of Household Savings and Labour Supply", *Economic Modelling*, Vol. 28, pages 2054-2070.

Waldron, M and Zampolli, F (2010), "Household Debt, House Prices and Consumption in the United Kingdom: a Quantitative Theoretical Analysis", *Bank of England Working Paper*, No. 379.