



BANK OF ENGLAND

Speech

The labour market, productivity and inflation

Speech given by

Martin Weale, External Member of the Monetary Policy Committee

At the Manchester Economics Seminar, Manchester

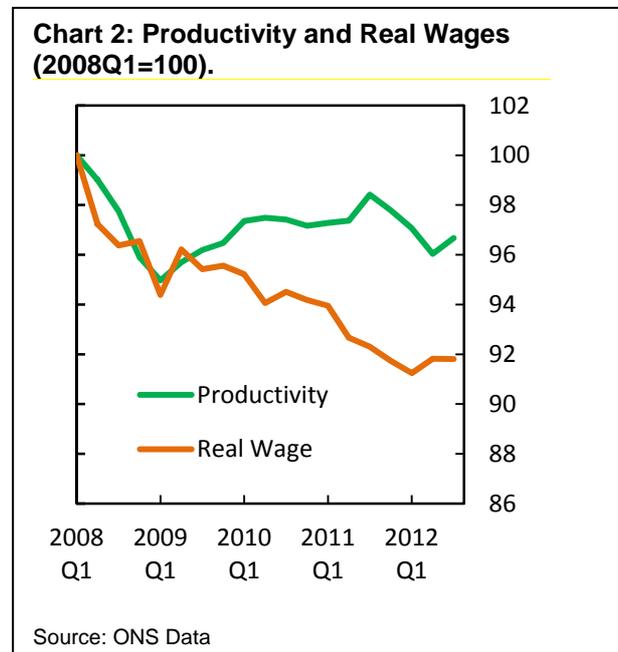
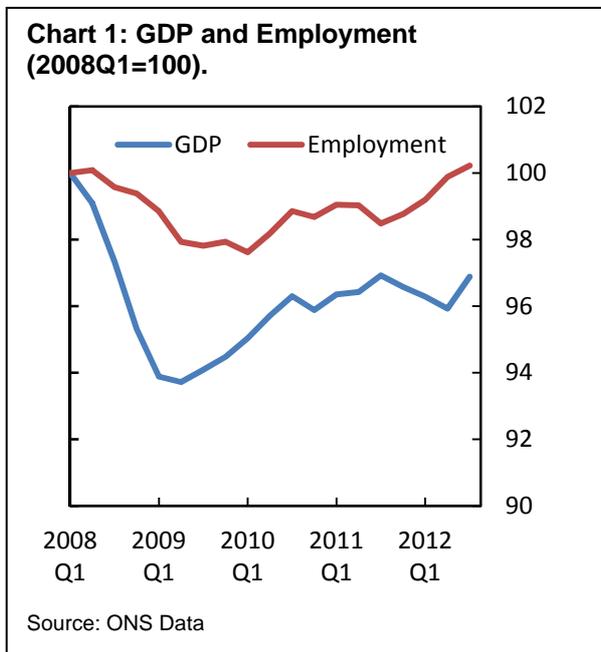
21 November 2012

I would like to thank Robert Gilhooly, Tomas Key and Tomasz Wieladek and for research assistance. 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.

Introduction

Thank you very much for inviting me to give this lecture. I would like to talk about the labour market and its recent performance. I will focus on the period since the economic crisis which affected us and the other advanced economies in 2008, but obviously I need some reference point, and the topic can be given more of a focus by asking whether and how far circumstances in the labour market have changed since the start of the crisis.

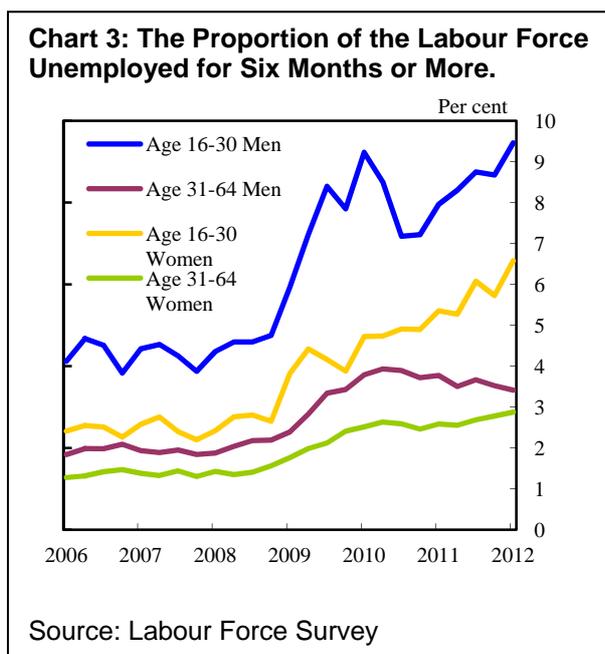
The last few years have, in terms of real GDP, seen sharp contraction, weak recovery and then stagnation, with movements in GDP in the last year largely explained by erratic factors. But, it is well known that, if one looks at employment data instead of GDP data, the British economy would appear to have recovered from the crisis. Chart 1 shows this. Although the unemployment rate is, at 7.9% markedly higher than it was before the economic crisis, overall employment was higher in the summer than it had been at the start of the recession in the first quarter of 2008. Of course the consequence of this is that productivity performance has been poor, as Chart 2 shows. This is an important issue for the Monetary Policy Committee because it seems unlikely that there can be a return to anything like normal growth in demand and output while productivity stagnates. Real wages have been even weaker than productivity; they have also been affected by changes to indirect taxes and by movements in the prices of traded goods in the aftermath of the depreciation of sterling in 2008.



As always, aggregates can mask a wealth of interesting and important detail. In this lecture I would like to peer behind the aggregates and show how different people have been affected differently in the aftermath of the crisis, I would like to begin with a look in more detail at the characteristics of the labour market, indicating what has changed since the crisis of 2008. I would then like to examine whether, and how far, labour market developments help explain what has happened to productivity over the last four and a half years. Finally I will examine the role of unemployment on wage and price inflation; an understanding of this is, of course, at the core of the Monetary Policy Committee's job of keeping the inflation rate close to two per cent per annum.

Patterns of Unemployment

Chart 3 shows unemployment of six months or more, calculated as proportions of the labour force. We can see from the charts that, for both men and women, much the sharpest change in employment opportunities has been for young people. I made the point earlier this year that, while the past few years have not been kind to many of us, they have been particularly unkind to those in the early years of adult life (Weale, 2012); this chart illustrates that.



Instead of looking at people of working age classified by age, I can classify them by qualifications. Charts 4 and 5 show overall unemployment rates on this basis, and here we see what I suppose I regard as the traditional pattern found in most advanced economies. Both men and women with good qualifications are less likely to be unemployed than are those with poor qualifications and the unemployment risks for the former have remained fairly low during the down-turn. Similarly redundancy rates plotted by qualification suggest that being well-qualified offers protection against redundancy risk in good times and in bad.

Chart 4: Unemployment Rates by Qualifications, Men.

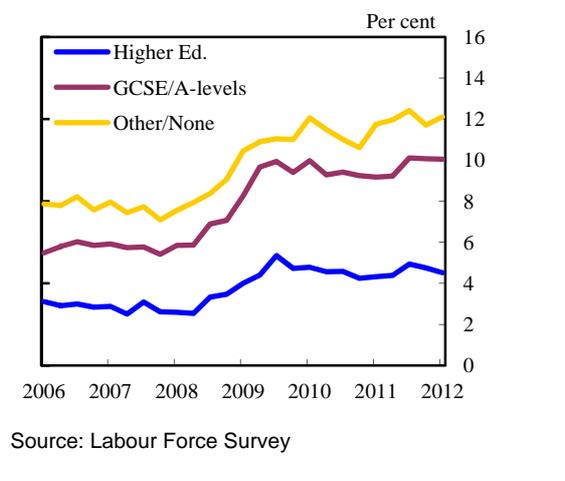
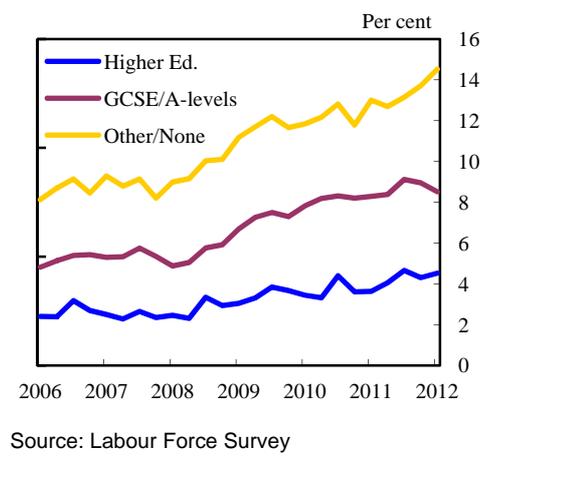


Chart 5: Unemployment Rates by Qualifications, Women.

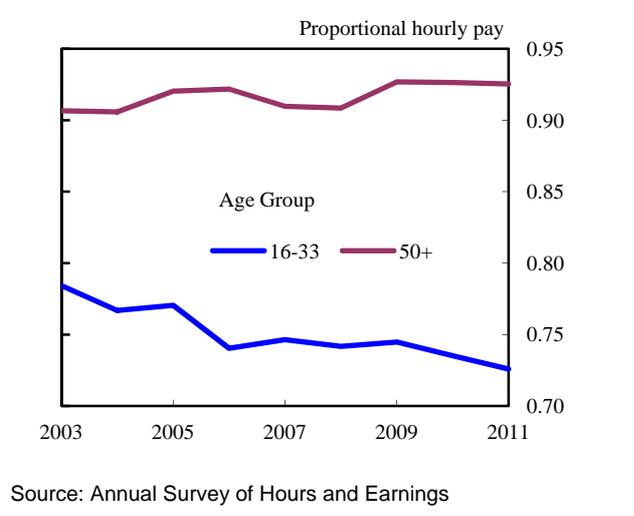


These charts remind us that, even though the overall labour market picture is reasonably favourable, unemployment is a very real problem for young people and for those with few qualifications.

Movements in Wages

The labour market is about wages as well as about employment and hours worked. I would also like to talk about wage rates. It is all too clear to many of us that real wages have fallen since the start of the economic crisis. But looking at data on hourly wages from the Annual Survey of Hours and Employment (ASHE), I can see how the wages of people in different age groups have changed. We know that, over the recent past real wages have moved downwards; I would like to focus on changes in the relative positions of people in different age groups. In Chart 6 I show wages of people aged 16-33 and those aged 50-64 as a fraction of

Chart 6: Hourly rates of pay as a proportion of the hourly earnings of people aged 34-49.



the earnings of those in the middle age group (34-49). The graph shows that the wages of both young and old workers are lower than those of people in the middle group. This is not a surprise. For many years it has been known that people's earnings tend to reach a peak, at least in relative terms, when they are in the forties (Mincer, 1974). But you can also see in this graph how the relative position of young workers has declined while that of the old age group has improved slightly relative to those who aged 34-49 since the middle years of the last decade and, to a smaller extent since the economic crisis. These data, together with the

earlier figures on unemployment, should bring home that the last few years have been very difficult for young people. No one should expect to remain unscathed while GDP remains some three per cent lower than at the start of 2008, as Chart 1 shows. But when people close to retirement draw attention to the effects of monetary policy on annuity rates, it is important to remember the economic reality faced by young people as summarised in data such as those of Charts 3 and 6.

Let me move on to discuss how people's wage rates have changed from one year to the next. Short-term aggregate movements in wages are conveniently measured by Average Weekly Earnings, the series the ONS computes from a monthly survey of businesses and used to compute the data in Chart 2. These data are core to our regular analysis on the Monetary Policy Committee. But they do not provide a deep insight into the functioning of the labour market. For that I again have to turn to ASHE. It is possible to follow people from one year to the next, so we can observe how individual earnings are changing.

An analysis of this by age appears, at first sight, to contradict my earlier observations. Chart 7 shows that, relative to the whole population, the wages of people in the young group have been rising fairly sharply¹, while those in the other groups have not kept up with the population average. But these data, unlike those of

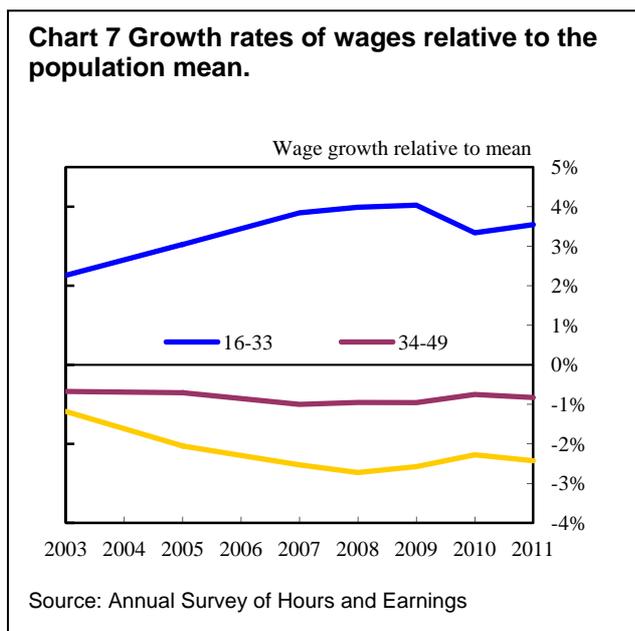


Chart 6 do not reflect the fact that, in each year, people move in and out of each population group; they simply reflect the pattern identified by Mincer (1974) that individuals' real wages rise in the early part of their careers, stagnate in the middle of working life and then may decline as retirement approaches. Young people on average experience wage rises as they age through the relevant age group but each year the thirty-four year olds migrate into middle age and are replaced by people just starting work, and earning generally appreciably less. So, taking Charts 6 and 7 together suggests that career progress has not changed very much but that the starting point for young people has fallen relative to that of the older age groups.

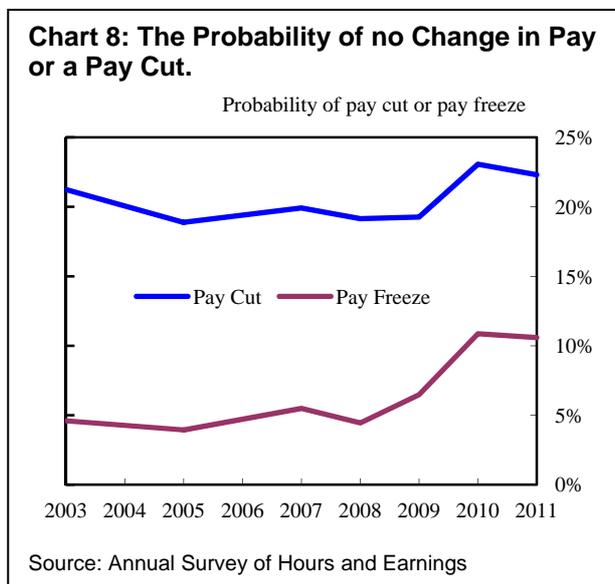
A general assumption in economics is that people's pay measures their contribution to output². This means that movements in real pay can be assumed to reflect movements in productivity and suggests that, the productivity of young people is not increasing relative to the rest of the population, as rapidly as it used to.

¹ Changes in ASHE mean that it is not possible to produce comparable data for 2004 or 2006.

² This proposition relies on the assumption that markets work reasonably well on average. It is obviously possible to think of individual examples which cast doubt on the idea that it is universally true.

While this undoubtedly is a depressing effect on productivity growth, it turns out that it accounts for a negligible amount of the decline since 2008.

The data also allow us to investigate whether sticky wages might be a factor impeding the adjustment of the economy. Chart 8 shows that, in the period before the crisis, it was moderately common for people to receive



no change in hourly pay³, and that reductions in pay were more likely than pay freezes, a point first observed by Nickell and Quetini (2003). The crisis itself had little effect on the proportions in either of these categories. But in 2010 and 2011 (not 2009 despite the fact that the period of sharp recession ended in the first part of 2009) both pay freezes and pay cuts became more common, with the increased proportion of people experiencing freezes being particularly marked. Dickens *et al.* (2007) suggest looking at the ratio of people experiencing freezes relative to those who experience freezes or cuts, as a measure of wage rigidity. This has risen from an average of under 20% before the crisis to just over 30% in 2010 and 2011,

suggesting increased rigidity despite the fact already noted that pay cuts have become somewhat more common.

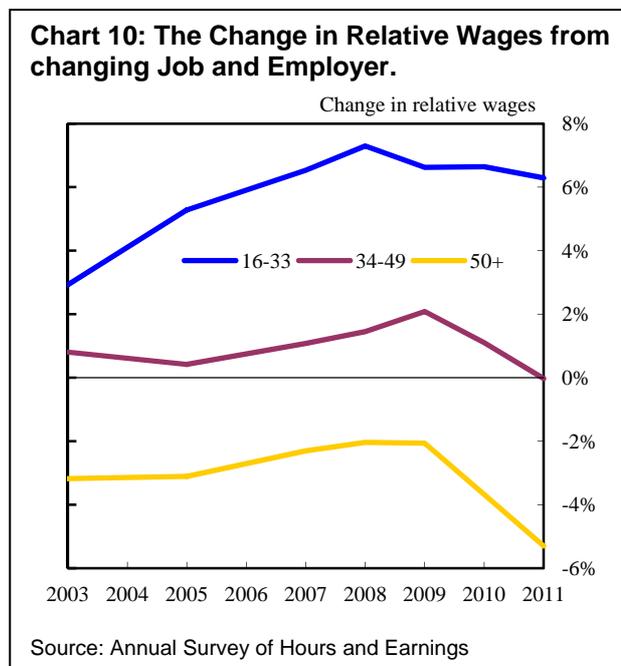
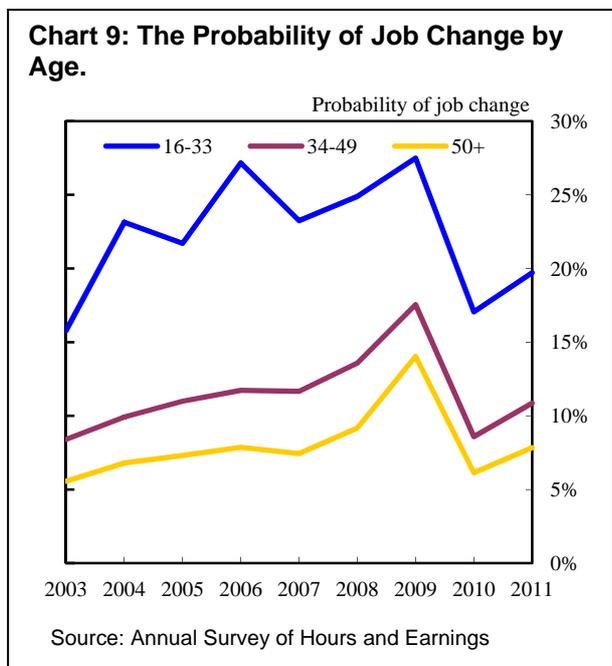
There is, nevertheless, an obvious difference between what these data show and the fact that the Bank's database of pay settlements shows very few pay settlements involved reductions in pay while, as we see, over 20% of employees experienced cuts to hourly basic pay. The apparent discrepancy between the settlement data and the individual data can be resolved only by matching the two, so that the experience of people working for specific employers can be related to what is observed on settlements. But it is in any case important to remember that settlements and *ex post* outcomes will always differ.

Job Dynamics and Productivity

We often hear about the importance of labour market flexibility. But how often do people change jobs and why? Data from the ASHE provide some interesting insight into this. Chart 9 shows the probability of someone changing job as a function of age. Not surprisingly, young people are the rolling stones, with a probability of changing job much higher than those of the older people. But do they really gather no moss? Chart 10 shows the growth in earnings for people who change jobs and change employer, measured relative

³ Defined as basic hourly pay moving by 0.2% or less over the twelve months to April in the year in question. Overtime payments and overtime hours are excluded from the analysis. The analysis here covers all people employed whereas Nickell and Quetini looked only at those who did not change their jobs and worked in the private sector.

to the average rate of growth of hourly earnings for the same group. An obvious conclusion is that young people change jobs to take advantage of new opportunities, or to move from less suitable to more suitable jobs, while old people may have job change forced on them as an alternative to retirement. But the chart also shows that, except for the youngest group, the benefits of changing job have declined in recent years; for people fifty and over the penalty has increased.

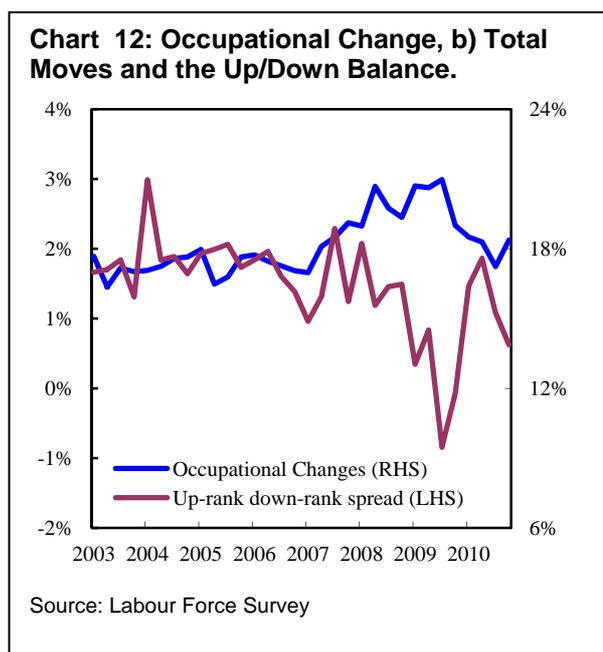
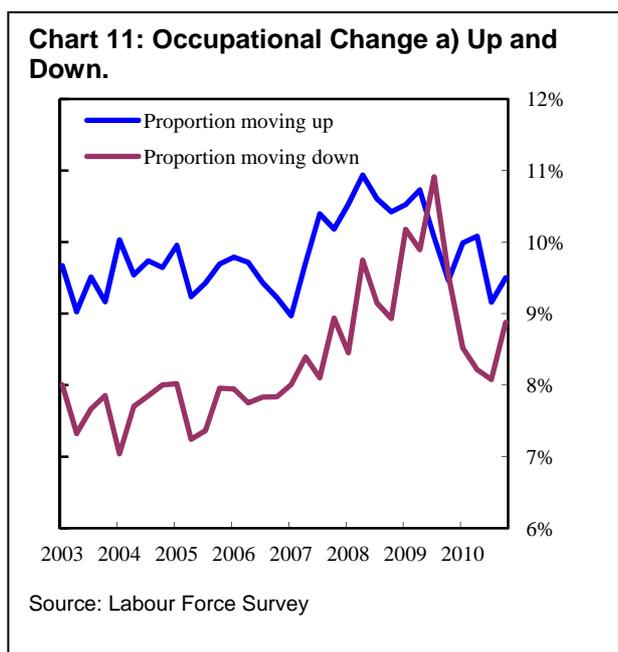


If I look beneath the surface, rather more seems to have changed. The Labour Force Survey provides information on the dynamics of the labour market and allows us to glean information not only on who is employed and who is not but also on what is happening to the people who are employed. To do this I rank occupations by their average earnings over the period 2003-2006 and look at people who move up or down the occupational ladder⁴.

Charts 11 and 12 show that between 2002 and 2006 there was a process of marked net *embourgeoisement*. People were appreciably more likely to move up the occupational ladder than down it. This may in part have reflected the labour market equivalent of grade inflation. Anyone who has run a business knows that it is costless to change someone's job title so as to enhance their self-esteem and, quite possibly their work satisfaction. But, I think it is also a natural state of affairs. As people develop their careers, they are likely to move to higher-status jobs. This does not mean that the workforce as a whole is necessarily upgrading because, as people retire from high-status jobs, they may be replaced in the workforce by young people

⁴ Classification changes mean that it is not possible to produce data for 2011.

working in low status jobs. In much the same way, there may be general upward drift in wages for much of



working life without the average wage necessarily drifting upwards.

In 2007 *déclassement sociale* became more common. The balance of downward moves relative to upward moves seems to have declined. This is, of course, of general interest. But it also can be used to extract information about people's productivity.

To the extent that these moves represent the evolution of people's productive power, or their human capital, and thus their impact on growth in labour productivity, that suggests that there may have been a change in the working of, or at least the opportunities offered by the labour market, with implications for the growth of labour productivity. So, following the work by Pratap and Quintin (2011) on Mexico, I would like to present a view of the effects of changing occupation on wages.

As you might imagine such an exercise is subject to all sorts of statistical concerns. Are the people who change occupation, whom I refer to as movers, comparable with those who do not move? If the answer to this is no then, it is not satisfactory to rely simply on a straight comparison of the two. The statistical means of addressing this issue is to fit a model which explains the probability that someone moves as a function of a wide range of their personal characteristics. I can then compare the wage growth of someone who does move with an average of their near neighbours, those who have a very similar probability of moving but in fact stay put⁵.

⁵ This technique is referred to as propensity score matching. See Dehejia and Wahba (2002).

The Labour Force Survey interviews respondents five times, at approximately quarterly intervals. It collects information on wages only at the first and last interviews. It has provided detailed information on occupational classification only from 2002 onwards, but this information is collected in every quarter⁶. There is some evidence of mis-coding; I have assumed people who change occupation or industry during their participation in the survey but end up where they started have simply been miscoded. But, after the adjustment for miscoding, the proportion of people who move during the course of the year averages 17.5% before the crisis and 19.4% afterwards. This compares with a figure of around one-fifth found for the United States (again after correction for miscoding; Kambourov and Manovskii, 2008) and Denmark (Groes et al., 2010). ASHE does not provide sufficient auxiliary information to make it possible to carry out a similar exercise. At the same time we have to remember that the rate of moves of occupation can be calculated only for those respondents to the Labour Force Survey who provide suitable information, and this is under one half of the total number of employed respondents to the survey. So more work needs to be done to establish whether and how far that part of the population for which we can do the calculations is representative of the employed population.

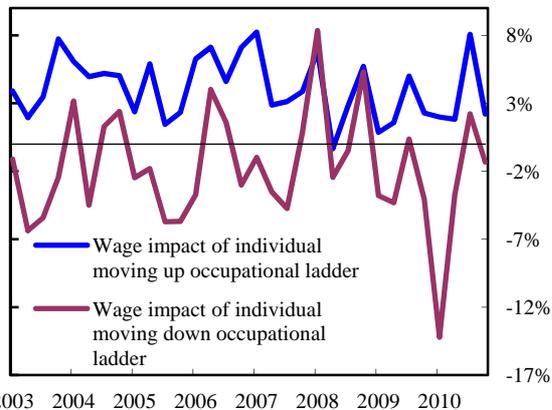
For the period for which calculations are possible, Chart 13 illustrates the estimated effect of moving up and down the occupational ladder on wage growth over four quarters. This shows clearly that while wage growth associated with moving in both directions was affected during the economic crisis, the impact of moving down the occupational ladder was substantially larger, with an average maximum wage penalty of fourteen percent, a non-trivial effect by any measure.

A notable feature of Chart 13 is that the wage impact of moving down the occupational ladder reaches its peak impact in 2010Q1, approximately two years after the onset of the recession in the UK. A delay of about two quarters can be accounted for by the fact that the Labour Force Survey is backward-looking. My estimates relate to people who changed their occupations in the year to the date in question. So on average they relate to occupational changes two quarters earlier than the dates shown. Chart 14 shows year on year movements in labour productivity together with the wage impact of moving occupation. There is little correlation between them.

But inspection of the chart suggest that the correlation would be substantially enhanced if the wage impact assumed to lag measured movements in productivity by about two years; this is consistent with the point made above about 2010 Q1. There is a good explanation why there should be a substantial lag. In the immediate aftermath of a fall in demand firms are likely to hold on to their labour forces in the hope that better times will return soon.

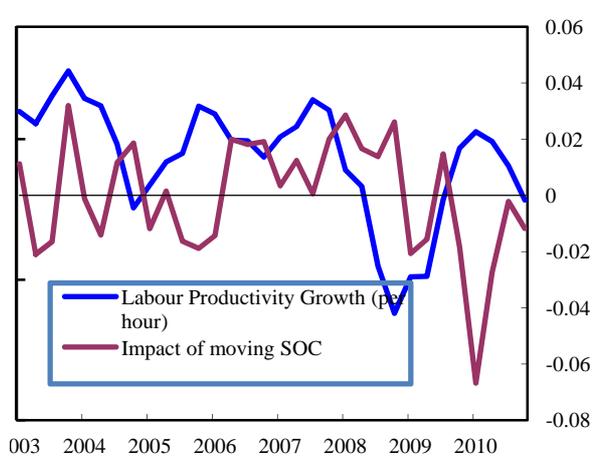
⁶ Once again it is not possible to produce estimates for 2011.

Chart 13: Wage Impact of Moving Up and Down the Occupational Ladder



Source: Labour Force Survey

Chart 14: Labour Productivity Growth and the Effects of Occupational Change



So output per hour worked will decline because labour is hoarded. Later on, as businesses become more realistic, people may find that they have to move occupation, taking what is available, in order to remain employed. Perhaps a sommelier is redeployed behind a bar.

As I mentioned, there is a question how far the people identified as moving occupation in the Labour Force Survey represent the full population. But if they were representative, the impact of the reduction in the moving premium from during the crisis would be to depress the rate of growth of real earnings by 0.5% p.a. After allowing for the fact that employee compensation accounts for about 60% of value added, this translates into a reduction in labour productivity growth of about 0.3% per annum. The total shortfall in labour productivity growth since the start of the crisis has amounted to about 3% p.a., so what I interpret as reduced labour market opportunity might account for just under one tenth of the shortfall. These calculations are carried out by comparing the period from 2008Q1 to 2010Q4 with the two years before the crisis. The figure is smaller figure if one compares the post crisis period with the whole of the period back to 2003. This is, of course, subject to further work on the question of whether adjustments are needed to deal with the matter of whether the data used are representative.

The *Inflation Report* has offered a number of other possible factors behind the productivity shortfall, but quantifying the effects of some of these has not proven to be as straightforward as the mechanisms I have described here. It is perfectly possible, of course, that, at the end of the day and putting all the components together, we will find that we have over-accounted, or under-accounted for the shortfall. In that case the various components can be adjusted with reference to their reliability as described by Stone (1980).

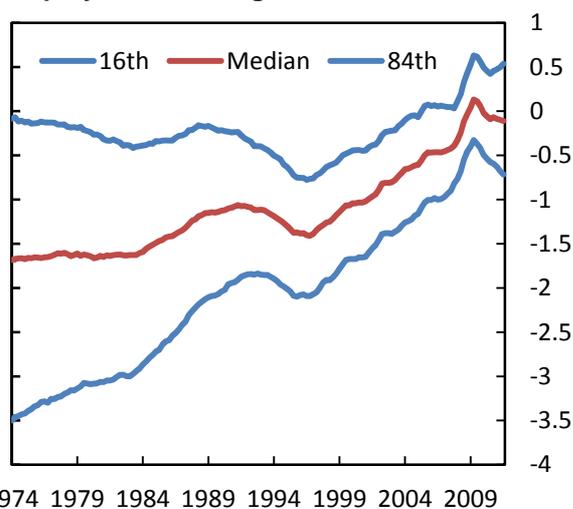
Unemployment and Inflation

So far I have presented to you a number of detailed indicators of the state of the labour market and the way in which it has changed. But from the perspective of someone whose job it is to deliver an inflation rate of two per cent, I also need to look at the relationship between labour market pressures and inflation. It is assumed, in the New Keynesian economic model that the inflation rate depends on the degree of spare capacity in the economy. In the simple model the degree of spare capacity is represented by the so-called output gap, the gap between what output would be were capacity fully used, and what output actually is. This New Keynesian Phillips curve contrasts with the original Phillips curve which sought to explain wage growth in terms of the unemployment rate. In the standard New Keynesian Phillips curve (Gali and Gertler, 1999; Woodford, 2003) there is no place for the direct effects of unemployment. More recently, however, Gali (2011) suggests that there is mileage to be gained by looking at the relationship between wages and unemployment.

Chart 8, which showed increased clustering of pay movements at zero, might be seen as a reason for expecting the relationship between the state of the economy and overall wage inflation to have weakened. But here I would like to explore that further, using an atheoretical approach to examine the drivers of inflation in the United Kingdom and how they have changed over the years. First I investigate the relationship between nominal wage inflation, unemployment and the interest rate. The framework is one in which the parameters of the relationship are allowed to evolve steadily over time (see Primiceri, 2005), without imposing any view on the way in which they change. No explicit allowance is made for inflationary expectations or the non-accelerating inflation rate of unemployment. But the time-varying nature of the system means that the effects of changes in these variables over time are readily absorbed into the time-varying constant term of the relationship explaining wages. The time-varying long-run coefficient for the effect of unemployment on wage inflation is shown in Chart 15.

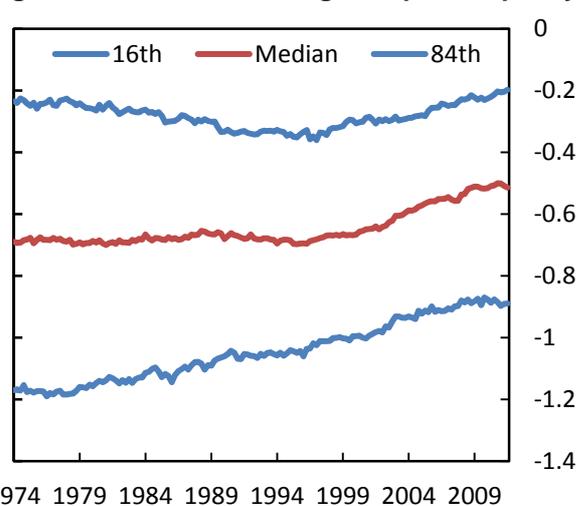
Seen from this perspective the impact of unemployment on wage inflation appears to have declined markedly. A policy-maker might well infer from this graph that unemployment has little effect on wage growth and that it would therefore be possible to add to demand, bringing unemployment down without any perceptible increase in wage pressure. Even then, however, there might be some grounds for caution because the graph does not show that unemployment has negligible effect on wages. Rather it suggests that there is about a 50/50 chance that the impact is zero or positive.

Chart 15: The Estimated Long-run Influence of Unemployment on Wage Inflation.



Source: Bank calculations

Chart 16: The Influence of Unemployment on Wage Inflation after allowing for Spare Capacity.



Source: Bank Calculations

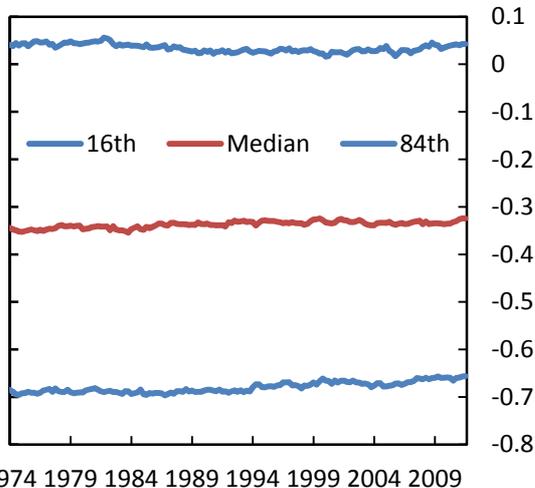
Should policy be set on the basis of what this graph suggests is marginally more likely or should be set on a basis which the graph does not exclude, the chance that unemployment depresses wages is just under 50%, given that there are good theoretical reasons to expect unemployment to depress wages at least to some extent? These are the sorts of questions that the MPC has to grapple with all the time. It does not have the luxury of making decisions only when the evidence is clear-cut.

But there is a broader issue. It is perfectly possible that the rate of wage growth is influenced by other variables in addition to unemployment and the interest rate. Spare capacity may also play a role. If employees know that the firms in which they work have spare capacity, they may be reluctant to push for wage increases. An analysis which allows for this possibility ought to provide a better picture if spare capacity is relevant, while leaving it unaffected if it is not relevant⁷. Once spare capacity is included in the system, a rather different picture emerges, as Chart 16 shows.

But even this might seem to be answering the wrong question because the MPC has the job of controlling inflation, not the job of controlling the growth of wages. Chart 17 shows the same analysis applied to RPI inflation (the CPI series is too short for this analysis). This now gives no hint of an upward drift. At the same time it should be mentioned that the chance that the impact of unemployment on inflation is zero or positive is put at something like twenty per cent- appreciably above that deduced from Chart 16.

⁷ This is always, of course, subject to the risk that the inclusion of too many variables may lead to over-fitting, or spurious explanation of the data.

Chart 17: The Influence of Unemployment on RPI Inflation.



Source: Bank calculations

You might think that these three graphs present a rather bewildering picture. The answer you get to the important question, what is the impact of unemployment on inflation, depends on the way in which the question is asked. But I think actually they serve to clarify the issue a great deal. It seems to me appreciably more likely than not that the current elevated rate of unemployment is playing a role in limiting wage growth and thus labour costs. At the same time I recognise, of course, that these charts do not, nor could they, prove this to be the case. There are undoubtedly other ways of looking at the data which will give a different conclusion and I can understand why some of my colleagues might take that view.

If unemployment continues to exert downward pressure on both wage and price inflation, why then should they both seem becalmed, the former at about two per cent per annum, and the latter somewhat above our target of two per cent per annum. As I mentioned earlier, the flexibility apparent in individuals' wages makes it unlikely that wage rigidity is the main factor responsible for the stability of aggregate wage growth. But one possible explanation is that the effect of unemployment is offset by upward pressure on real wages as a response to the recent squeeze on living standards. Anyway, with stagnant or falling productivity, wage growth of even two per cent per annum is at best barely compatible with the inflation target.

Other Inflationary Pressures

A feature of the New Keynesian framework, used by the Bank in its medium term analysis, is that there is an element of rigidity in the system because price and wage adjustments are infrequent. Wages and prices are, nevertheless, ultimately sensitive to demand. But our experience over the last few years has been that administered prices can have a separate influence on inflation; this may continue for some time to come. I am not convinced that the New Keynesian model describes the behaviour of an electricity or water company which sets its prices through negotiation with a regulator. Or of the universities whose increase in tuition fees has added 0.3 per cent to inflation and with further contributions in October 2013 and October 2014. In addition to these effects, inflation is, as my colleague, Paul Fisher, pointed out (Fisher 2012) also affected by the behaviour of industries which relate price increases to past movements in the Retail Price Index rather than to their own circumstances.

Of course the rate of inflation is always an average. Some prices rise faster than inflation and others fall or rise more slowly. But it is quite an act of faith to believe that other price-setters will put up their prices less because they anticipate spending power being reduced as a result of these administered price changes, at least without the Bank of England maintaining a tighter monetary policy than would otherwise be the case.

While inflation remains above target, there is an obvious risk that, at some point, people involved in setting of wages and prices will start to think that we do not take the inflation target serious. It need not be a sudden change. The risk is all the greater because the change is more likely to be very gradual; I am reminded of the point that a frog placed in a pan of water which is gradually heated up will eventually die instead of jumping out. Movements in the gilts market suggest that expectations of RPI inflation have, if anything moved down in the last few months. But signals from financial markets are not necessarily relevant to people's perceptions in the rest of the economy. We obviously cannot rely indefinitely on the belief that independence of policy setting is enough to maintain credibility.

The report by David Stockton (2012) reminded us, the MPC has repeatedly forecast inflation falling below target but that outturns have been higher. The inflation figures last month produced another unwelcome surprise, although the broader picture is that inflation has fallen much as we forecast last year. But, in the light of the pressures mentioned above, I think it is more likely than not that inflation will remain above target for much of the next two years. My analysis suggests that additional stimulus would, without any corresponding improvement in productivity, add to inflation.

There is, nevertheless, an argument for a further stimulus. It is possible as David Miles (2012) has argued, that a revival of demand would lead to a sharp improvement in productivity growth from sources other than those discussed here. Indeed some improvement in productivity growth is already built in to our recent forecast. But, at the moment I do not feel we have a quantitative understanding of the factors contributing to weak productivity clear enough to be confident that productivity would move in line with a sharp increase in demand.

Conclusions

In this lecture I have wanted to convey three very important points about the state of the economy. First, the labour market situation, as it has evolved over the last few years has been particularly difficult for young people. Indeed, while the consequences of low annuity rates for people retiring have received considerable attention, I suspect that, overall, young people have fared worse than those approaching old age. Secondly, there seems to have been a change in the working of the labour market which has resulted in there being fewer opportunities for career advancement through changing occupation or industry of employment than there were in the years before the crisis. This might account for up to ten per cent of the shortfall in productivity, but there remain questions about how representative the data are before it is possible to come

to a firm conclusion. Thirdly, despite some evidence of increased wage clustering around zero growth, I do not find clear evidence to support the view that the effect of unemployment on inflation has declined markedly since the crisis.

The implications of these last two points are that sustained above-target inflation remains a concern. While it might be true that a sharp increase in demand would lead to an improvement in the productivity position, this is not a certainty. Such a policy would need to be justified only on the basis of a balance of risks and not on the basis of a sure outcome.

References

- Dehejia, R.H. and S. Wahba. (2002). "Propensity Score Matching Methods for Non-experimental Causal Studies". *Review of Economics and Statistics*. Vol 84. Pp 151-161.
- Dickens, W.T, L. Goette, E.L. Groshen. S. Holden, J. Messina, M.E. Schweitzer, J. Turunen and M.E.Ward. (2007). "How Wages Change: Micro Evidence from the International Wage Flexibility Project". *Journal of Economic Perspectives*. Vol 21. Pp 195-214.
- Fisher P. (2012). Interview in *The Sun*. 27th September.
- Gali, J. and M. Gertler. (1999). "Inflation Dynamics: a Structural Econometric Analysis". *Journal of Monetary Economics*. Vol 44. Pp 195-222.
- Gali, J. (2011) "The Return of the Wage Phillips Curve". *Journal of the European Economic Association*. Vol. 9, Pp. 436-61.
- Groes, F., P. Kircher and I. Manovskii (2010). "The U-shapes of Occupational Mobility", Manuscript, University of Pennsylvania.
- Kambourov, G., and I. Manovskii (2008). "Rising Occupational and Industry Mobility in the United States: 1968-1997", *International Economic Review*, 49(1), Pp. 41-79.
- Miles, D. (2012). Monetary Policy and a Damaged Economy. Speech at the Society of Business Economists Annual Conference. 24th May
- Mincer, J. (1974). *Schooling, Experience and Earnings*. Columbia University Press. New York.
- Nickell, S and G. Quintini, 2003. "Nominal wage rigidity and the rate of inflation". *Economic Journal*. Vol. 113(490), Pp 762-781.
- Pratap, S. and E. Quintin (2011). "Financial Crises and Labour Market Turbulence". *Journal of Monetary Economics*. Vol 58. Pp. 601-615.
- Primiceri G. (2005). "Time Varying Structural Vector Autoregressions and Monetary Policy". *Review of Economic Studies*, Vol. 72, Pp. 821-852
- Stockton, D. (2012). *The Monetary Policy Committee's Forecasting Capability*. Presented to the Court of the Bank of England.
- Stone, J.R. N. (1980). "Whittling away at the Residual". *Journal of Economic Literature*. Vol XVIII. Pp. 1539-43
- Weale, M. (2012). From retailers' paradise to shoppers' strike: what lies behind the weakness in consumption? Dean's lecture, Cass Business School. 29th February
- Woodford, M. (2003). *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press.