Ben Broadbent: Compositional shifts in the labour market

Speech by Mr Ben Broadbent, Deputy Governor for Monetary Policy of the Bank of England, at the "Understanding the Great Recession: from micro to macro" Conference, Bank of England, London, 23 September 2015.

* * *

Accompanying charts and table can be found at the end of the speech.

I would like to thank Will Abel and Rebecca Burnham for research assistance and I am 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.

I want to begin by thanking the organisers for putting together this excellent conference and for asking me to come along. I'm honoured to be in such distinguished company. Up front, I also want to thank Bank economists – especially Matt Corder and Alina Barnett – for their help. If this were a working paper it would have multiple authors.

Today's tagline – "from micro to macro" – is a sobering one for macroeconomists. We're used to describing the economy in broad, sweeping terms, using numbers on aggregate output, total employment, average wages.

Yet, in practice, the modern economy is hugely diverse and, underneath the surface, highly dynamic. New ideas replace old, businesses fail and are formed and, even when the net change in employment is small, and the labour market in aggregate looks stable, the gross flows in the labour market can be huge.

This fact, highlighted many years ago by Steve Davis and John Haltiwanger in the US, was no less apparent during the great recession in the UK. Employment fell by over half a million between 2008 and 2010. But even over that period there were over nine million new jobs created, over ten million if you include job-to-job flows. On average, over the past 10 years, gross inflows and outflows into and out of employment have run at an annualised rate of around one sixth of the stock (Chart 1).

As I say, this has been long understood. Nor should it be forgotten that, whatever the complicated texture underneath, the aggregate economic data still exhibit significant empirical regularities over the business cycle. Chart 2 is one example. Using annual data during the inflation targeting period, the black line plots annual growth in average earnings, less inflation expectations, against the rate of unemployment. At least until the past couple of years, when AWE growth appeared unusually weak, the fit was pretty good. Higher unemployment has gone hand-in-hand with lower real wage growth. Uncannily, the regression line was pretty much exactly the same as that identified by A.W. Phillips, using data from the classical gold standard period (1870–1913), in the 1958 paper that gave the relationship its name (that regression line is in blue).

Yet, from time to time, what happens under the surface has material effects and can distort these aggregate relationships. Today I want to describe one such example, involving compositional shifts in the labour market. In recent years, these have been significant: specifically, both the jobs lost in the 2008–09 downturn, and those regained during the recovery, have been skewed towards the lower-paid end of the spectrum. This has imparted a bias to average pay and productivity growth – upwards in the recession, downwards in the past couple of years – over that period.

Chart 3 plots these estimated "compositional effects" on average wage growth against the change in unemployment. They went up during the recession, boosting average earnings growth, and fell back sharply in the recovery, when unemployment declined. Last year, on these estimates, the skew towards low-paid job creation knocked 0.3% off average wages, relative to those of the median. This compares with a pre-recession average of +0.7% per year.

One key point I want to make today is that, directionally at least, this isn't at all unusual. It's true everywhere that employment tends to be more cyclical for the young, the less experienced and those with fewer qualifications. Chart 4, for example, illustrates a similar phenomenon in the United States: those without degrees experience much larger swings in employment than college-educated people. As a result, measures of average earnings growth in the US are biased upwards in recessions (relative to the experience of people continuously in employment) and downwards in recoveries. That's why many people pay more attention to the Employment Cost Index[1] of US wages, which adjusts for these compositional effects, than to the (unadjusted) headline earnings series.

What is a bit different about the recent UK experience, however, is the scale of the effects. Compared with the cycle of the 1990s, they appear to have been less positive during the recession – in the more recent downturn job losses weren't as skewed towards the lower paid – but more negative during the recovery (the skew in net job creation has been bigger). Overall, taking the 2008–14 period as a whole, the compositional effects added around 3% less to average wages than in the equivalent six years in the early 1990s. If anything, Chart 3 hints at a break in the series long before the Great Recession, sometime in the late 1990s.

It's not clear why this occurred (I'll later offer some slightly speculative suggestions). But the recent drag from compositional effects might help to explain a couple of recent puzzles in the aggregate data.

One is low productivity growth. As we know, output per employee has been particularly weak in recent years. It fell steeply during the recession and its rate of growth has been pretty insipid since. Chart 5, plotting three-year growth rates for output and employment, is one way of expressing that. What this work on compositional effects suggests – somewhat tentatively, and only if you're prepared to interpret it in this rather crude fashion – is that this may partly reflect a slower rate of improvement in the average "quality" of the employed workforce. "Human capital", to use the economists' terminology, has been growing at a slower rate, particularly compared with the first half of the 1990s.

Note the word "partly". I don't think this can really be seen as the main, or even a major, part of the productivity puzzle. Even if you're prepared to view these compositional shifts as some independent phenomenon, unrelated to other things that have depressed productivity growth – and that is questionable – their scale isn't really big enough to make that much of a dent. That 3% figure, which is probably a ceiling on the true contribution[2], compares with a productivity shortfall of over 10% since the recession.

The results offer clearer help with a second puzzle, namely the apparent ease with which real average earnings have adapted to lower productivity growth. The two must match each other in the end[3]. But, in the past, it's been hard for the labour market to absorb negative

---

[1] The Employment Cost Index (ECI) measures the change in the cost of labour by excluding the influence of employment shifts among occupations and industries. More information about the ECI can be found here In the UK, the ONS has an experimental series – the "Quality-Adjusted Labour Input" – which makes a similar adjustment: the latest release can be found here.

[2] Some of the factors that might have reduced productivity growth might also have affected the skill mix of employment, so it's not clear whether this can legitimately be seen as independent effect: in the limit, saying productivity growth is weak "because" the economy has created more low productivity jobs has a slightly circular flavour to it. Note too that, at least in the short run, a given compositional effect on wages should not be translated one-for-one into productivity space – you'd expect the latter to be smaller (I explain this point in footnote 10). So the 3% figure should be viewed as a ceiling for the estimated impact of these shifts on productivity over the 2008–14 period.

[3] Real wages can diverge from productivity either because average profit margins are changing or because the relative price of what firms produce diverges from the price of what consumers buy. In an open economy, for example, an improvement in the terms of trade will raise real incomes relative to productivity; the upward trend in profit margins in the US has depressed wages relative to national income. There has been no such trend in

shocks to real take-home pay – whether from lower productivity, higher taxes or rises in import prices – without the unfortunate devices of higher inflation and/or higher unemployment. In other words, such shocks tend to put upward pressure on the NAIRU (at least initially, they move you rightwards along the Phillips curve in Chart 2, not downwards[4]). Yet, for much of the past two years, pay growth was weak even as unemployment declined sharply. I commented on the striking flexibility of real wages in a speech last summer, from which Chart 2 is taken.

Suppose, however, that the drop in productivity has had nothing to do with existing jobs but instead reflects the skew in the creation of new jobs towards relatively less productive, less well-paid roles. To that extent, there's no real-wage resistance, among existing employees, to overcome, and the arithmetic impacts on average pay and average productivity are necessarily immediate and coincident.

This would also help to explain why pay surveys, which track wages of the median employee (rather than the average), appeared to over-predict wage growth in 2013 and 2014 (Chart 6): they failed to take account of changes in the mix of employment. You can see from Chart 6 that when the compositional effects were positive, during the recession, the surveys under-measured AWE growth.

Finally, these shifts can also shed light on the most recent trends in the aggregate data. The latest estimates show that the compositional effects have risen during the first half of this year – the skew towards lower-paid jobs has diminished. This might account for the acceleration we've seen in average productivity and wages; it could also explain why there's been a narrowing of the gap between AWE growth and the surveys.

So having summarised the main points let me describe in a little more detail where the data comes from and how these effects have been estimated. I'll then say something, if only tentatively, about why they might have differed from the last cycle – why, in particular, the compositional effects have been more negative during the recovery, and perhaps before that as well (my hands will be waving slightly at that point). A short final section concludes with some remarks about policy.

Generating the estimates

To understand the results it helps to know something about where the data come from and how they're constructed. The first and most important point is that the headline measure of wages in the UK, the AWE, is what it says it is – an average. Employers are asked simply to divide their total wage bill by the number of employees. So if – say – a responding firm were to substitute a less experienced, less well-paid person for someone better paid, its submission to the AWE would fall even if no individual's wage had changed. It is these "compositional effects" that we are trying to measure and account for.

We do so using data from a different source – the Labour Force Survey – which is sent out to households. The main questions in the LFS are about employment status: do you have a job and, if not, are you looking for one? (If your answer to the second question is "yes" you are counted as unemployed, if "no", economically "inactive".) As well as these headline numbers,

---

the UK and the terms of trade have been broadly constant in recent years. Thus the only source of real wage growth, over time, is productivity growth.

[4] The inverse relationship between the NAIRU and productivity growth has been noted in many countries (Staiger et al (2001), Slacacek (2005)) nor is real wage resistance the only possible reason for the link. More micro-founded approaches to the labour market, based on the costs of job search, also predict a negative relationship between trend productivity growth and the natural rate of unemployment. In that setting, in which employers trade off the cost of search against the prospective surplus from a new job, lower expected growth in productivity reduces the present value of that surplus and lowers the rate of job creation. The canonical description of this "capitalisation" effect is due to Pissarides (2000).

the ONS's monthly labour market release publishes standard summary statistics about the duration of unemployment and the reasons for inactivity.

But, at least since 1994, the LFS has also included a host of questions about the characteristics of employees and their jobs. These are listed in Table 1. Some are easy enough to answer – your age, where in the country you work and long you've been in your job. Others, those related to respondents' qualifications, or the sector they work in, offer a long list of potential answers. There are over eighty sectors from which to choose, hundreds of detailed occupational descriptions. Table 1

People are also asked about their pay. And, unsurprisingly, there are clear correlations between wages and characteristics. Pay is higher in some sectors and occupations than others; it's also higher, at the level of the individual, for people who have better qualifications and more experience, in terms both of age and job tenure. Overall, if you take the post-1994 sample as a whole, the spread of individuals' characteristics statistically explains around one half of the variation in their pay.

This distribution can also change over time. Individuals' characteristics can evolve – people become more experienced, in principle they can also switch occupations or become better educated. The distribution will also shift if there are differences between the characteristics of those entering and leaving employment. We've already seen how large those gross flows can be. So even a relatively modest contrast between the two groups can lead to significant changes at the aggregate level. The compositional effects, which are simply the fitted values of a regression of pay on characteristics, attempt to track the effects of these shifts from one quarter to the next. In aggregate, the estimated effects amount to a "hedonic" adjustment for labour quality. Chart 7, summarising the results, is a version of a graph we published in the last two Inflation Reports.

Note that, over time, the effects have tended to be positive. The average age and experience of the workforce has grown; more people have gone into higher education and higher-paid occupations.

But on top of that underlying trend, we've seen sharp movements, in both directions, during the recent cycle. Both the drop in job creation during the Great Recession, and the corresponding pick-up during the recovery, were skewed towards the less skilled, less well paid end of the jobs market. Chart 8, which plots the differences in job creation across three classes of occupational characteristics, depicts one aspect of that pattern[5]. You can see similar differences in rates of job creation between young and old, less and more educated. And as I outlined in the introduction, the effect of these swings was to add materially to the growth of average earnings during the downturn and to subtract from it in 2013 and 2014.

Unskilled employment generally more cyclical

Before asking why this has happened let me make a few points about the estimates.

First, there are some approximations involved. One involves the mix of surveys: we're estimating the impact of compositional changes on LFS pay data and assuming the effect simply carries over to the AWE. That's the most reasonable assumption, but it's an assumption nonetheless. Also, because of the way we've done this regression – for the sake of simplicity it's been estimated in one go, using the entire 20-year sample – we're implicitly

[5] As I said, there are potentially hundreds of detailed occupational descriptions from which LFS respondents can choose. The ONS's standard scheme then makes high-level classifications – "high skill" occupations, for example, whose job creation rates are plotted in blue in Chart 7, are defined as "managers, directors, professional, associate professional and technical"; "low skill" refers to sales and "customer services, process and plant operatives, elementary".

assuming that changes in each of the characteristics have constant effects over time[6]. But there's a substantial economic literature explicitly devoted to explaining *changes* in some of these coefficients, notably the labour-market returns to better qualifications (the "skill premium"). I'll come back to this point later.

Second, the estimated compositional effects have been driven mainly by changes in the mix of people, and their individual characteristics, not by shifts in employment across sectors. Pay clearly differs across industries, even controlling for those characteristics. And on occasion sectoral shifts can matter for average earnings. In Chart 7 the marginal contribution of changes in employment across industries is in red. Thanks mainly to disproportionately large job losses in the financial sector, where pay is higher than average, these were negative in 2003 and 2009. But in the rest of the sample, they don't do much: the compositional effects have been predominantly *within*, not across, sectors. During 2014, for example, it's the fact that new jobs went to people with below-average qualifications and experience, and in lower-paid occupations, that depressed average pay growth[7].

Third, as I emphasised in the introduction, this cyclical pattern is entirely typical. The fact that new hires have low job tenure – and that the average experience of the workforce therefore declines when unemployment falls rapidly – is hardly surprising. That average skill levels should be counter-cyclical is, perhaps, less obvious. But it is a widely and long observed empirical fact: it's true in all countries that employment rates are lower and more cyclical among the young and unskilled.

Why should low skilled employment be more cyclical? Economists have come up with a number of possible explanations. One likely factor is that skilled jobs tend to be more specialised. That means they're also more likely to be retained, or "hoarded", in a downturn: the costlier it is to replace someone the more reluctant will be the employer to let him or her go.

Another possibility is that low-skilled workers are more easily substituted by new capital and new technologies. So when firms are relatively thinly capitalised, early on in an economic expansion, they tend to hire more unskilled people. Later on, as the recovery matures and investment in new capacity starts to take hold, it's the demand for skilled employment that starts to grow faster.

These are both things that would tend to boost the relative *demand* for high-skilled labour during a downturn. Differences in labour supply might also play a part. For example, some have argued that those with low wages have less to lose from exiting employment. Their pay is lower relative to unemployment benefit, especially during downturns. They may also have fewer options: when vacancies are scarce, skilled people have the capacity to move into lower-paid jobs while searching for new opportunities (Nickell (1979), Chassamboulli (2011)).

Why have compositional effects been less positive during this cycle?

These factors can help to explain why employment is more variable for the low paid. It's possible the same factors have been at work in the most recent cycle in the UK. But the last

---

[6] The regression also includes time dummies so allows for arbitrary variation of pay growth net of the compositional effects.

[7] In principle, demand-side factors could still matter. Perhaps, for example, the problems of the financial industry, whose employees are better qualified than average, have pushed employment towards other industries. To that extent, the regression could be picking up effects whose origins are demand-led – a shift in output across sectors – but attributing them to changes in individuals' characteristics. But this seems unlikely: in truth, shifts in sectoral employment haven't been big enough to generate such effects. Even if you exclude individual characteristics from our regression – the shifts in qualifications and occupations that do most of the empirical work – variations in sectoral employment shares don't do a good statistical job of explaining wage growth.

few years do look distinct, at least compared with the early 1990s cycle – overall, the compositional effects have been adding less to wage growth. They were less positive during the downturn, but have been more marked, in a negative direction, during the recovery. Overall, the "quality" of the workforce looks to have been growing at a somewhat slower rate since the onset of the 2008 recession – worth around ½% point a year on average wage growth – than it did during the equivalent period in the 1990s.

Why might this have occurred? I can think of two possibilities, though there are no doubt others as well. One is that the financial crisis has depressed the relative demand for high-skilled employees. I mentioned earlier that economists believe high-skilled jobs are probably complements, rather than substitutes, for capital and new computer technology. This is the most common explanation for the marked rise in the "skill premium", in the US and elsewhere, during the 1980s and 1990s: new technologies favoured skilled jobs and their wages rose faster than those of the unskilled[8]. So if the incentives to invest in new physical capital have been depressed since the Great Recession, by things like higher rise premia or problems in the banking system, perhaps the same applies to investment in human capital.

It's also possible the relative supply of low-skilled labour has risen. We know, for example, that migration flows have risen significantly over the past fifteen years, and that immigrants are more likely to take low-skilled jobs than UK-born workers (see, for example, the Report of the Migration Advisory Committee (2014)). So if easier immigration has made the overall supply of labour more responsive to economic conditions in this country – if, as if often argued, it's flattened the Phillips curve (Bean (2006), Bentolila (2008)) – it's probably done so to a greater extent for low-skilled than high-skilled workers. So when the UK economy grows faster than its neighbours, as has been the case over the past couple of years, you'd expect to see greater inward migration and a disproportionate rise in the supply of low-skilled labour in particular.

I came across a possible example of this on a recent Agency visit. The head of a small business told me his firm had been looking to hire an accountant. They'd originally wanted someone very experienced who might in time fill the shoes of the finance director. However, they soon realised that the only way they could get someone like that was to bid them from an existing job within the UK – apparently a costly and lengthy process. So they decided instead to employ two much more junior and less well trained people. One worked for them already; the other, from outside the UK, they found in relatively short order after advertising on an EU-wide basis. The firm would end up being less productive than it might have been (two people are now doing the work of one). But there'd be a compensating fall in its average wage and, given the relative availability of less experienced employees, it made sense for them to go down this route.

What I'm suggesting has more to do with relative quantities, not relative wages, and more to do with the short than the longer run: at a time when the UK has been outgrowing its neighbours, you'd expect easier immigration to amplify the rise in the relative supply of low-skilled employment.

Some concluding remarks

Let me sum up and, in doing so, say something about the implications for monetary policy.

As macroeconomists, monetary policymakers like to deal in simple aggregates and averages. Output is doing this, employment that and wages the other. We know that the

[8] More recently, economists have argued that computer technology substitutes most directly for the low-to-medium skill tasks. This would help to explain why the changes in the distribution of pay have had a "U" shape, with positive effects not just on high-skill pay but that at the low end as well (Goos and Manning (2003), Autor et al. (2006)).

experiences of individual firms and households are extremely diverse but like to think we can nonetheless lump them all together and treat them as homogenous wholes.

Much of the time that's a reasonable thing to do. But sometimes it matters what's happening underneath. As we've seen, there was a distinct skew in recent job creation towards the less well paid in 2013 and 2014, one that has subsided in the first half of this year. While not the only factor – narrowing slack in the labour market has probably been more important – this helps to explain why growth of average wages was weak in those two years and has since picked up a bit.

At least at the margin, these shifts will have had similar effects on productivity. So, in and of themselves, they have no direct bearing on impact on inflationary pressure[9]. Indeed, because the compositional effects are highly counter-cyclical, and tend to go up and down with unemployment, perhaps it doesn't matter much if you never took them into account and stuck with Phillips curves estimated on unadjusted wage data.

The problem is that, although they're qualitatively similar in different cycles, the compositional effects do appear to vary in scale. This means it's worth allowing for them: changes in the mix of employment may not affect inflationary pressure but they can have an impact on the aggregates we normally use to assess that pressure. Thus, in general, our estimated Phillips curves fit better when you use AWE growth adjusted for compositional effects. In the past two to three years in particular, the extent of the skew towards low skilled employment helps to explain some of the weakness of average productivity and pay.

That skew might reflect the particular circumstances of today's recovery. Perhaps easier immigration has made low-skilled labour easier to come by, especially at a time when other European economies haven't been growing as quickly as the UK. Maybe deterrents to investment in physical capital, and new technologies, have also reduced the relative demand for high-skilled labour. If so, then you might expect the compositional effects to improve further over time: a strengthening of the recovery in the rest of Europe would reduce the relative supply of low-skilled labour and a fall in risk premia on new investment, in human as well as physical capital, would raise the demand for high skills.

I guess we shall see.

Thank you!

References

Autor, D., L. Katz and M. Kearney, 2006, "The polarisation of the U.S. labour market", *NBER* Working Paper no. 11986.

Bean, C., 2006, "Globalisation and Inflation", Bank of England Quarterly Bulletin, 2006 Q4.

Bentolila, S., D, Juan J. and J. F. Jimeno, 2008, "Does immigration affect the Phillips curve? Some evidence for Spain" European Economic Review, Elsevier, vol. 52(8), pages 1398-1423, November.

[9] For the economists amongst you let me be more a little more precise: one would expect an impact on average unit labour costs but not on the marginal costs that matter for inflation. Suppose workers are paid wages equal to their marginal revenue products. In that case, all else equal, an extra job will have no impact on the ratio of the two (marginal costs). It will also have the same impact on aggregate output, in levels terms, as on aggregate wages. But because output is bigger than wages – GDP also includes gross capital income – the proportionate effect on output will be smaller, by a factor equal to the capital share of income (around a third). So when the compositional effects are positive – say they add 1% to average wages – you'd expect a slightly smaller proportionate rise in average productivity (around 2/3%) and a small rise, of around 1/3%, in the ratio of wages to productivity (average unit labour costs).

Bernanke, Ben S. and Powell, L. James, 1986, "The Cyclical Behavior of Industrial Labor Markets: A Comparison of the Prewar and Postwar Ears." in: R.J. Gordon, ed., The American Business Cycle, Chicago: *University of Chicago Press*

Chassamboulli, A., 2011, "Cyclical Upgrading of Labor and Employment Differences across Skill Groups", B E Journal of Macroeconomics

Goos, M. and A. Manning, 2007, "Lousy and lovely jobs: the rising polarisation of work in Britain", *The Review of Economics and Statistics*

Nickell, S., 1979. "Education and Lifetime Patterns of Unemployment*," Journal of Political Economy, University of Chicago Press,* vol. 87(5), pages S117–31, October.

Phillips A. W.,1958, "The Relation Between Unemployment and the Rate of Change of Money Wage Rates in the United Kingdom, 1861–1957", *Economica.*

Pissarides, C., 2000, "Equilibrium Unemployment Theory", *Massachusetts Institute of Technology, 2nd edition*

Slacalek, J., 2005, "Productivity and the Natural Rate of Unemployment", *Department of Macro Analysis and Forecasting, DIW Berlin*

Solon, G, R. Barsky, and J.Parker*,* 1995, "Measuring the Cyclicality of Real Wages: How Important is Composition Bias?" *Q. J. E. 109 (1994), pp. 3–25.*

Staiger, D., J. H. Stock and M. W. Watson*,* 2001, "Prices, Wages and the U.S. NAIRU in the 1990s", *NBER Working Paper* no. 8320.

Chart 1: Large flows into and out of employment

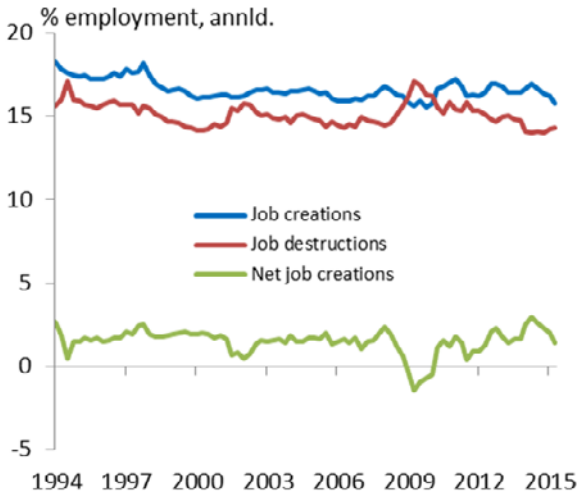

Source: Labour Force Survey and own calculations

Chart 2: Growth of average wages weak in 2013, 2014 despite low unemployment

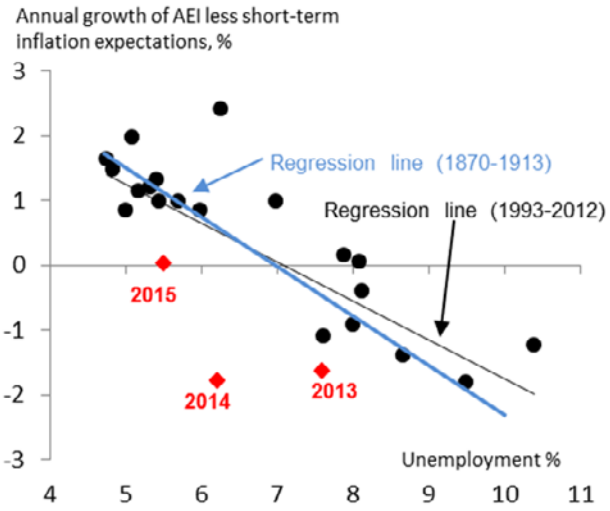

Source: ONS, Bank of England, NIESR and Hills, Thomas, Dimsdale (2010)

Note: AEI until 2010 and AWE since; inflation expectations as implied by indexed gilt market; 2014 diamond is for H1.

Chart 3: Compositional effects currently weak, generally counter-cyclical

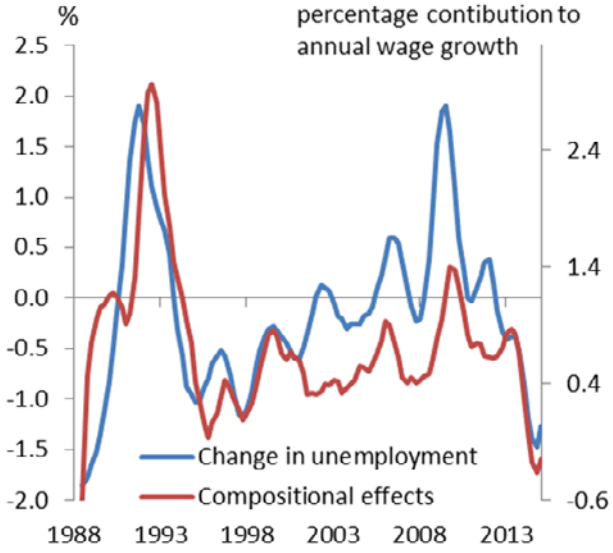

Source: ONS, Labour Force Survey and own calculations

Chart 4: Low skilled employment more cyclical in the US

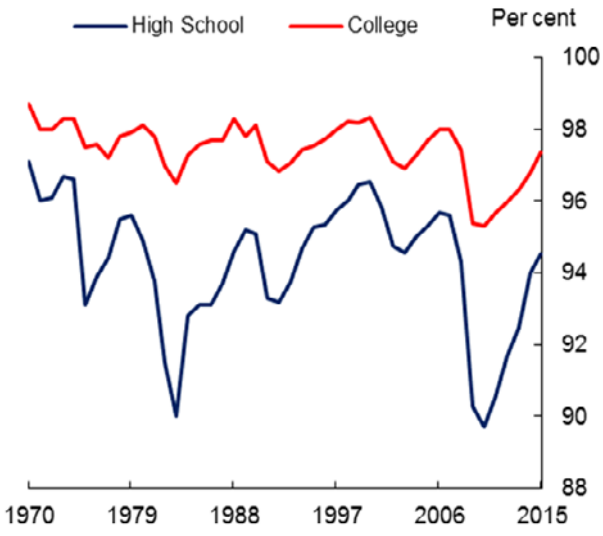

Source: U.S. Bureau of Labour Statistics

Chart 5: Productivity growth very weak in recent years

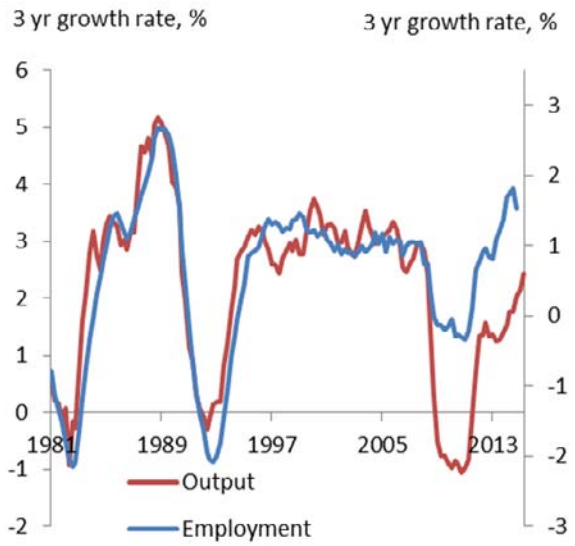

Source: ONS

Chart 6: Pay surveys better match for AWE ex compositional effects

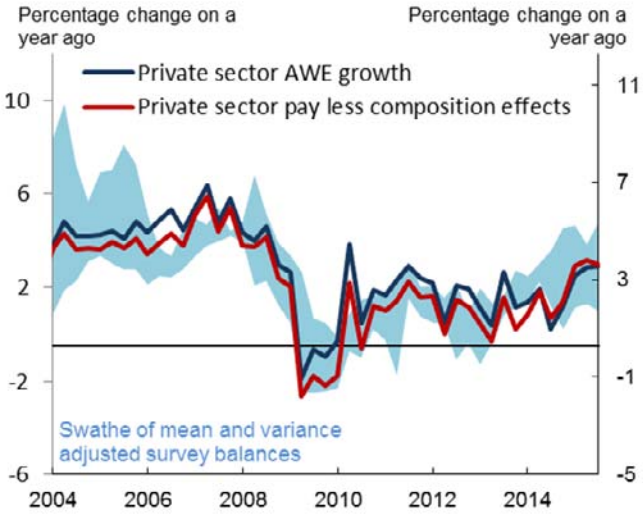

Source: ONS, CBI and BCC

Table 1: Detailed questions in the Labour Force Survey

Characteristics	LFS Question
Age	How old are you?
Occupation	What was your (main) job in the week ending Sunday the [date]?
Tenure	In which year did you start working continuously for your current employer? And which month was that?
Education	This next section is about qualifications gained in the UK. Please think about ALL qualifications you have ever gained [in the UK], even if it was a long time ago or you are not using them now (please exclude expired qualifications).
Sector/ Industry	What did the firm/organisation you worked for mainly make or do (at the place where you worked)? Enter a title for the industry

Source: Labour Force Survey (user guidance)

Chart 7: Compositional effects driven by characteristics of employees, not sectors

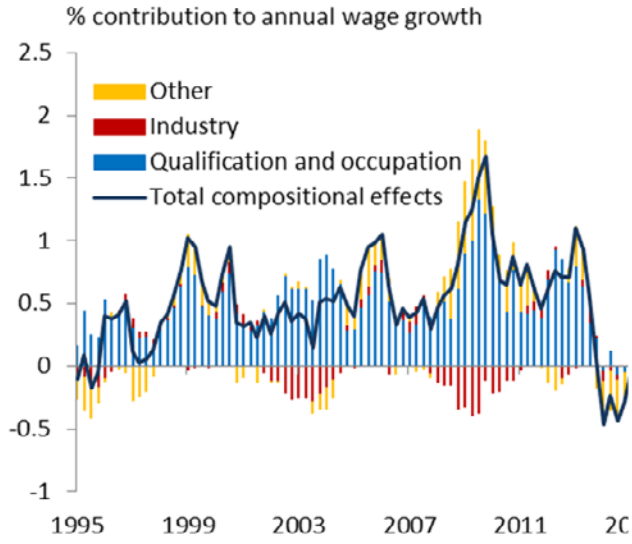

Source: Labour Force Survey

Chart 8: Job creation fell further, and has since recovered faster, for low skilled

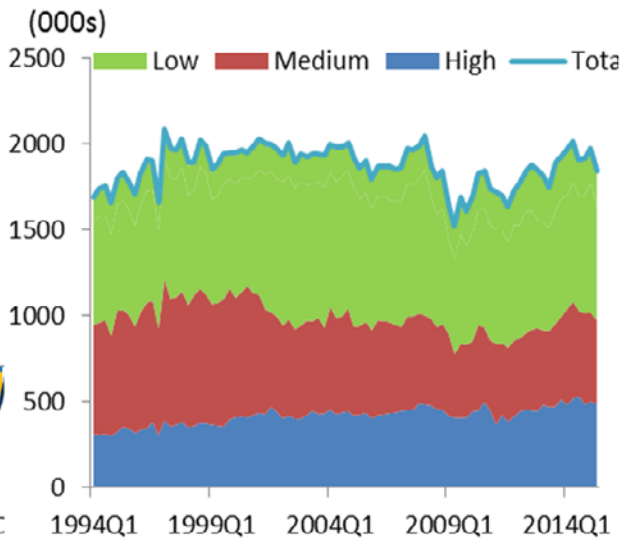

Source: Labour Force Survey

Chart 9: Over longer sample, shifts in average qualifications driven compositional effects

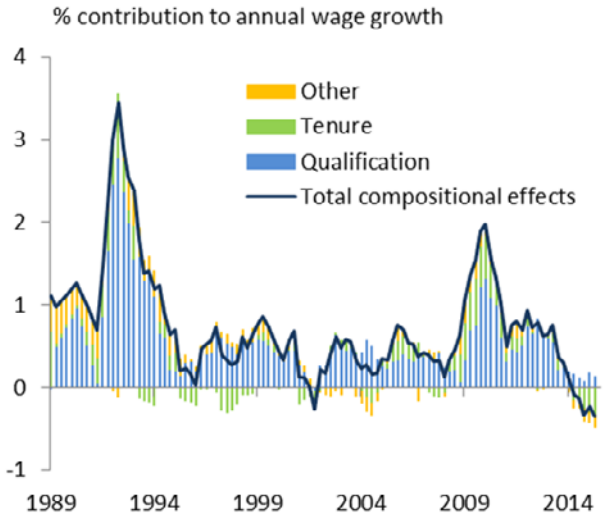

Source: Labour Force Survey

We can't demonstrate this with the detailed LFS dataset from 1994, as it includes only one proper cycle. But the positive relationship with unemployment exists even during the Great Stability (1994–2008). And, at the cost of a smaller set of variables – the survey didn't ask people about their occupation and sectors before 1994 – we've managed to extend the estimates back a few years, to the late 1980s. The aggregate series you've already seen in Chart 3[10]. The more detailed breakdown is in Chart 9. You can see clearly the same cyclical pattern during the last-but-one cycle: steeper falls in employment for the less well qualified during the recession, resulting in a positive compositional effect, the opposite pattern in the upswing.

[10] The longer series simply splices on the pre-1994 estimates, based on the narrower sample, to the existing post-1994 series. One might ask how similar the two sets of estimates are in the overlapping post-1994 period, and the answer is that they're close. This is partly because the information about sectors doesn't do much in explaining variations in pay growth (given the other characteristics). It's also because occupations and qualifications are well correlated, so when you've got only one of them it picks up quite a bit of the variation in the other.

Chart 10: Rise in employees with degree, slight fall in degree premium

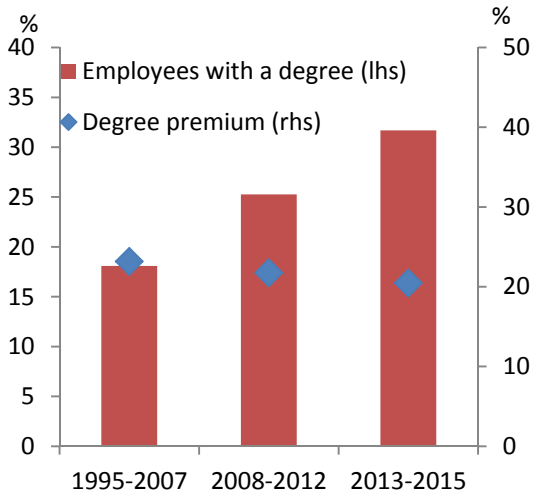

Source: ONS and Bank calculations

One shouldn't exaggerate the labour-market consequences of easier immigration. Whatever its impact on the supply of low-skilled labour in the short term, it's not been enough over time to have a noticeable impact on relative rates of pay. Chart 10 plots an estimate of the UK skill premium, specifically the marginal effect on wages of having a degree. We've derived this by running the same regression of pay on employee characteristics, but on a year-by-year basis rather than across the sample as a whole. Perhaps because the supply of graduates has steadily risen, the estimated degree premium has actually fallen slightly over the past twenty years[11].

[11] It's also important to note that, although immigrants are disproportionately represented in low-skilled occupations, they are not, on average, any less well qualified than UK-born people (to the extent domestic and foreign qualifications are comparable). In fact, according to the Migration Advisory Committee, one reason some employers have preferred immigrants for low-skilled jobs is that they "have higher level qualifications than their low-skill job requires".

Chart 11: Real pay growth more variable when adjusted for compositional effects

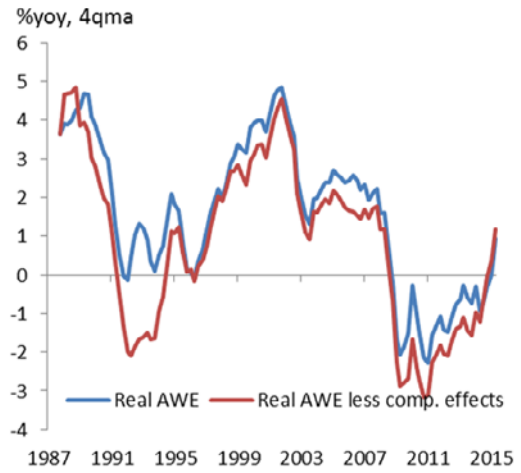

Source: ONS and Bank calculations

To be sure, you'd tend to under-estimate the cyclicality of "underlying" wage growth[12] and the true slope of the Phillips curve. Over the past few years, as in earlier cycles, real pay growth for people who stayed in their jobs[13] has been more variable than the average (Chart 11). But if it's a really regular thing – if the compositional effects always move in lock step with unemployment – then a Phillips curve estimated on an unadjusted AWE series would still take them into account, only implicitly so.

[12] This point was made by Bernanke (1987) and, in a well-known paper, by Solon et al (1995).

[13] To be precise, the red line is pay growth for someone with fixed characteristics – fixed educational qualifications and occupation, unchanging experience. Since any individual's experience is bound to grow over time (if only by virtue of getting older) there is no such person in real life! A more accurate description of the red line is average wage growth for people continuously in employment, adjusted for any the impact of changes in average age and experience across the workforce.