Wages, productivity and the changing composition of the UK workforce

By Will Abel, Rebecca Burnham and Matthew Corder of the Bank's Structural Economic Analysis Division.[1]

- Over the past 30 years the composition of UK employment has changed substantially — these changes have important implications for wage and productivity growth.
- These 'compositional effects' can be more prominent during times of increased labour market change and may have dragged down on wages over the past two years.
- The drag from compositional effects is likely to fade as the labour market normalises, pushing up on both productivity and wage growth.

Overview

Wage growth is a key indicator of inflationary pressure in the economy, and is important for the Monetary Policy Committee.

This article discusses how the changing characteristics of those in employment can affect aggregate measures of wages. If a disproportionate amount of particularly low or high earners enter or leave employment in a given time period, these 'compositional effects' can alter the average level of wages and hence the measurement of wage growth. Such effects have been particularly large during the financial crisis and the ensuing recovery.

Since the mid-1990s the increasing quality of the labour force has on average added close to half a percentage point to wage growth per year (**summary chart**). This growth has been driven primarily by an increase in the average education level of the workforce and a shift towards high-skilled occupations. These shifts mean workers are more productive, which in turn feeds through into higher wages.

At the start of the recession, compositional effects pushed up even further on wage growth relative to normal, as lower-paid employees were laid off and struggled to find new jobs.

But in 2014 and 2015 these effects went into reverse, slowing wage growth relative to normal. At its peak the drag from compositional effects was estimated to be close to 1 percentage point. This was caused by both lower-skilled employees returning to the workforce, and also by higher-skilled employees exiting the workforce.

For most of 2015 this drag was holding down wage growth by around ¾ of a percentage point relative to normal, but in the most recent data this has started to dissipate. The drag is likely to continue to dissipate as the labour market normalises, leading to an increase in both wage and productivity growth.

Over the past 20 years, there appear to have been substantial changes to how some characteristics affect wages. For example, the estimated returns from having a degree have fallen, as have the returns from job tenure. The decline in relative wages for those over 50 compared with younger age groups, which has been typically observed in the UK labour market, has disappeared since the crisis. The estimates also suggest that regional and gender pay gaps have narrowed over recent decades.

Summary chart The contribution of compositional effects to average annual wage growth by time period

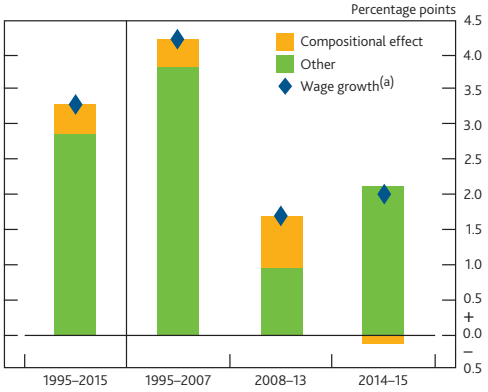

Sources: Labour Force Survey, ONS and Bank calculations.

(a) Whole-economy total average weekly earnings.

(1) The authors would like to thank Rosetta Dollman for her help in producing this article.

Wage growth has been weak during the UK economy's recovery from the financial crisis. An understanding of the reasons for this is important when assessing how much inflationary pressure there is in the economy, and thus the appropriate stance for monetary policy. One potential explanation for weak wage growth is that aggregate wage measures have been affected by the changing composition of the UK workforce during the recession and its aftermath.

This article explores the extent to which compositional effects can help to explain weak wage growth in recent years. It sets out how the structure of the UK workforce has changed over the past 30 years. It then explains how, against this backdrop, the cyclical effects of the changing composition of the labour force on wage growth can be estimated and applied to the recent recession. The final section looks at how the effect of some key characteristics on wage levels has changed over the past 20 years.(1)

Setting the scene: the persistent weakness of wage growth

Wage growth is typically thought to be driven by three main factors: productivity, inflation expectations and labour market slack (that is, the level of hours worked relative to potential labour supply). In the long run, an individual's real wage should track their productivity — the value of what they produce over a given period measured in real terms. Nominal wages also need to take account of inflation to preserve spending power. As wages are typically set in advance, this means factoring in inflation expectations when setting wages. Finally, in the short run at least, if unemployment is high then workers are more likely to accept lower wages because their alternative options, such as searching for a job elsewhere, are worse.

From 2001–07, annual wage growth averaged around 4¼% (**Chart 1**). This can largely be accounted for by a combination of steady productivity growth of around 2% and stable household inflation expectations. The unemployment rate was close to its estimated equilibrium rate over this period and is likely to have exerted little pressure on wage growth.

By contrast, from 2008–12, wage growth averaged less than 2%. Some of this fall can be explained by the sharp fall in productivity growth. But even accounting for this, wage growth was weak during the crisis and the immediate recovery. This is likely due to the high level of labour market slack, reflected in the rise in unemployment, reducing the pressure on employers to raise wages to retain and attract staff.

Wage growth from 2013–14 remained weak at around 1¼%. A rise in wage growth might have been expected as workers regained bargaining power due to declining slack in the labour market, but pay growth remained weak even though unemployment began to fall rapidly in 2013.

Chart 1 Unemployment and total wage growth 2001–15(a)

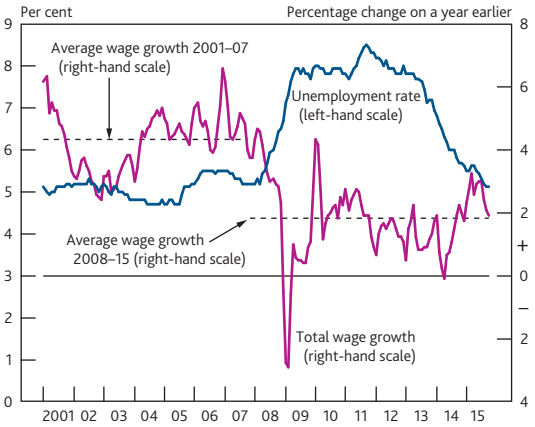

(a) Wage growth is measured as average weekly earnings.

Wage growth remained below its pre-crisis average in 2015 despite unemployment falling back towards its pre-crisis rate of 5%. Although wage growth picked up somewhat during 2015, it eased back to around 2% at the end of 2015.

The failure of wage growth to pick up as much as might have been expected following the fall in unemployment in recent years means that other factors could be affecting wage growth.

The potential role of changes in the composition of the workforce

One explanation is that 'compositional effects' have dragged down on measures of productivity and wage growth over this time.

Compositional effects in this context mean the change in aggregate measures of wage and productivity growth caused by changes in the mix of characteristics of those in employment. The timeliest measure of wages in the United Kingdom — average weekly earnings (AWE) — is calculated by dividing the total amount paid as wages in the economy by the total number of employee jobs. Similarly productivity per worker is simply the total value of all output produced in the United Kingdom divided by the number of workers. If some types of worker are more productive than others and therefore earn more than others, then as people enter and exit the workforce, switch industries, undergo training and gain experience, the resulting changes in composition will have an effect on the average level of both wages and productivity.

(1) These calculations only cover 20 years as opposed to the 30 years initially examined because of data limitations.

To illustrate this, in 2000 about a sixth of workers in the United Kingdom had a degree; by 2015 the proportion of workers with a degree had risen to about a third. Workers with a degree tend to earn more than others. Consider a simple example where workers with a degree earn £10 an hour and all other workers earn £7 an hour. In these conditions, a doubling in the proportion of workers with a degree would, all else equal, raise the average level of wages from £7.50 to £8 (**Figure 1**). This change in the average level of wages is solely the result of compositional effects; the pay of individual workers is unchanged.

Figure 1 Stylised example of effect of changes in the composition of the workforce

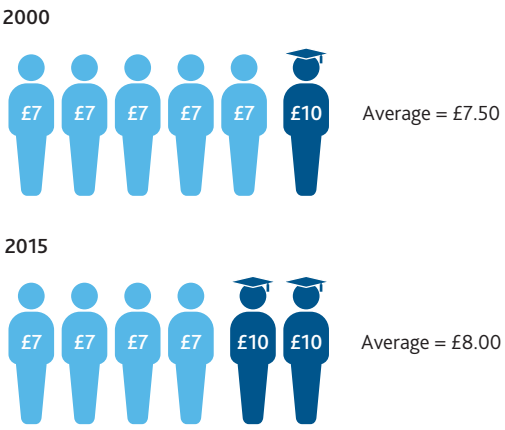

It might be expected that gradual changes in the composition of the workforce would not affect cyclical wage dynamics. Even during the rapid recovery in employment seen in 2013 and 2014, annual employment growth peaked at only 2.3% — a share of the workforce too small to drag up or down on the average wage growth of the entire employed population substantially.

But small changes in employment mask substantial churn in the labour market. The changes in aggregate employment are the product of much larger flows into and out of employment as well as the movement of people between different jobs. On average, just under 3½% of the workforce leave employment every quarter, with slightly more flowing in from either inactivity or unemployment (this slight gap reflects the increasing level of employment over time). A further 2%–3% of employed individuals move between jobs each quarter (**Chart 2**). The cumulative effect of these flows means that, on average, in each year between 1994 and 2007, around one in ten workers had changed employer compared to a year before and one in ten had entered the workforce from either inactivity or unemployment: this leaves about eight out of ten who were at the same employer as they were the previous year. Even without incorporating the changes caused by people being trained and moving roles within the same company, this level of churn can generate substantial changes in workforce composition.

Chart 2 Flows into, out of and within employment every quarter

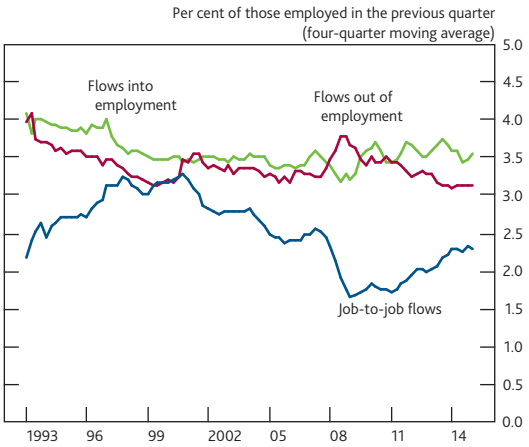

Source: Labour Force Survey.

While such flows and compositional changes are continuously occurring, the cyclical fluctuations in labour market flows that have been seen throughout the recession and its aftermath are likely to make these effects particularly pronounced. For example if, at the beginning of a recession, those who lose their jobs are disproportionately low-skilled workers, average wage measures would be higher than they would be otherwise, boosting measured wage growth. The compositional effect on wages here would be positive.

Understanding the size of compositional effects is important for monetary policy decision-making. To the extent that such cyclical shifts are only likely to affect wage growth temporarily (until the shifts in the mix of employment are complete) they may mask underlying shifts in pay pressures. Furthermore, wage increases are not always inflationary. If wages and productivity increase equally, then the wage cost per unit of output, known as the unit wage cost (UWC), will remain unchanged. It is changes in UWCs — rather than nominal wages — which should in theory be most important for companies' pricing decisions and hence inflationary pressure in the economy.(1) As compositional effects would be expected to affect both productivity and wage growth in a similar way they should not have much impact on UWCs or inflationary pressure. But an understanding of compositional effects helps the MPC interpret movements in pay and productivity alongside other important factors such as slack in the labour market.

Structural changes to the composition of the UK workforce

In order to analyse the compositional changes caused by cyclical behaviour during the financial crisis and subsequent

(1) Technically the primary concern is with unit labour costs, which include non-wage costs (such as pensions), rather than unit wage costs; for the purposes of this article this distinction can be ignored.

recovery, it is necessary to understand the substantial longer-term structural trends that have affected the composition of the UK labour force.

An individual's pay and productivity are affected by factors such as their level of education, work experience, the industry they are employed in and their specific job role. Across the economy as a whole, there have been large shifts in the distribution of most of these measures over recent decades.

Probably the most substantial change to the UK labour force, in terms of earnings potential, has been the rapid increase in the educational attainment of workers. In 1985, less than a tenth of workers had a degree: 30 years later, a third have a degree (**Chart 3**). As the workforce becomes more skilled over time, this should raise the average level of pay and productivity.

Chart 3 Highest education level of those employed[a]

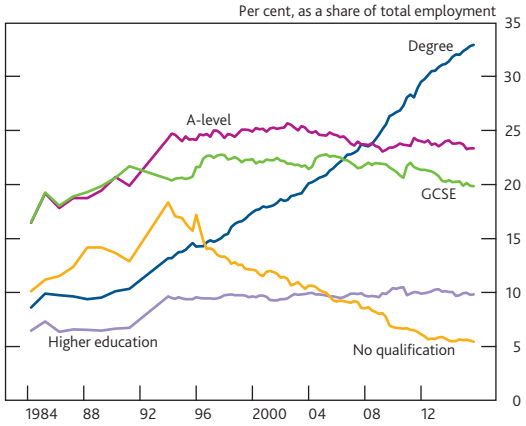

Sources: Labour Force Survey and Bank calculations.

(a) There are slight kinks in the time series due to changes in the survey methodology; where highest level of qualification has been recorded as 'don't know' these results are not shown.

As the supply of skills has increased in the economy so too has the share of high-skilled jobs. Determining the skill level of any specific job can be difficult, but using a commonly accepted approach where managers and those in professional or technical roles are considered highly skilled, the share of UK employment in such occupations has increased from 34% in 1992 to 44% in 2015.[1] Roles such as these are likely to be more productive and so be more highly paid.

The industries in which people work have also changed substantially — most notably due to the decline of manufacturing employment since the late 1970s and corresponding increase in the share of service industries in employment. Manufacturing accounted for a quarter of employment in 1978 compared with 8% today, with services now accounting for 83%. The effect of this shift upon earnings is, on the surface, ambiguous. Some of the increase in services employment has been in professional and technical industries, which are likely to be more highly paid. However the largest increases, in terms of employment, have come in the education and healthcare sectors, neither of which are particularly highly paid relative to other industries. The financial and information technology industries have not expanded substantially since the late 1970s, going from a combined share of employment of 6% to 7%. This contrasts with the fact that by other measures these industries have grown substantially (for example, stock market capitalisation).

Another important shift has been the ageing of the UK population (**Chart 4**). The ageing of 'baby boomers' and the decline in the UK birth rate has resulted in an older society, with 43% of the population aged 45 and over in 2014 compared with 37% in 1981. This shift has fed through into the workforce: over the same period, the share of those 45 and over in employment has risen from 36% to 43%. Older workers are likely to have greater levels of work experience and thus productivity than younger workers, although these effects might not be uniform — for example, productivity may decline when individuals near the end of their working lives.

Chart 4 The ageing of the UK population

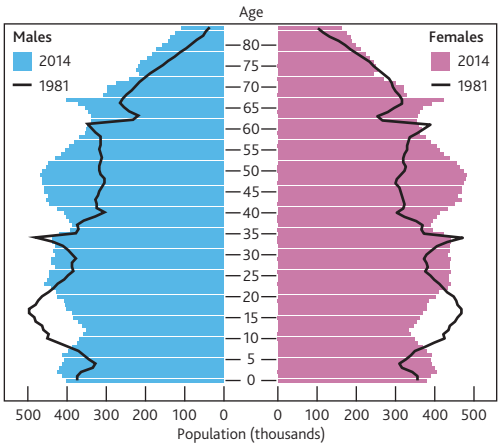

Female participation in the workforce has also increased over time. The share of women in employment has gone up from 41% in the mid-1980s to 47% in 2015. Numerous economic studies show that there is a gender pay gap faced by women[2] relative to men which is unexplained by socioeconomic factors like education and the industries in which they work. To this extent, women's increasing share of the workforce might have pulled down on overall average pay. While this effect on individual pay is statistically significant, the effect is likely to be substantially smaller than factors such as an individual's education, occupation and age.

(1) Measures of occupation are only taken back to 1992 due to changes in the occupation classifications in the Labour Force Survey.
(2) See, for example, Leaker (2008).

The share of people working part-time has also risen, from a fifth in the mid-1980s to 27% now. The rise partly reflects the rise in female participation mentioned above: while women accounted for 47% of all employment in 2015, they only accounted for 37% of full-time employment. Part-time employees are more likely to experience lower wages per hour compared with their full-time counterparts.

The importance of these compositional shifts for aggregate wages will depend on the difference in pay for different types of workers. **Chart 5** shows the average wage in 2015 for groups of individuals based on different factors, either the skill level of their occupation, their highest level of education or their gender.(1) The difference in wages between these groups, based on any one criterion alone, is substantial. This explains why changes in the share of employment in these different groups over time give rise to compositional effects on pay. However, these simple averages can capture a variety of inter-related effects. For example, those with higher levels of educational attainment are more likely to work in high-skilled occupations. To be able to understand fully the role of compositional effects, the marginal effect of different characteristics needs to be identified. This is discussed in the next section.

Chart 5 Average earnings in 2015 for different groups

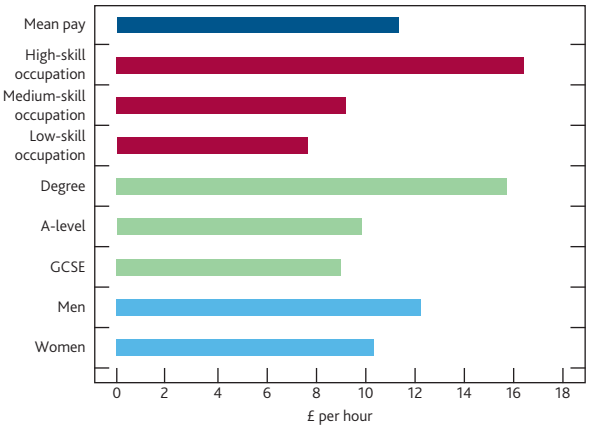

Source: Labour Force Survey.

Calculating compositional effects

To calculate the size of compositional effects on wage growth, estimates are made of how much different characteristics affect wage levels, which are then applied to the changes in the mix of employment over time.

In order to do this, quarterly anonymised individual-level data from the Labour Force Survey (LFS) are used, which includes a detailed set of individual and job-related characteristics, including age, gender, industry, occupation, job tenure and education as well as the self-reported level of pay. These data are available on a quarterly basis from 1994 onwards.(2)

By regressing an individual's observed hourly wage against a range of these individual and job-related characteristics, it is possible to calculate for any given quarter how these factors affect their wage, on average.(3) These wage equations allow for a decomposition of observed changes in wage growth into a combination of 'explained' and 'unexplained' variation. Here, the 'explained' variation is the measurement of compositional effects, where changes in pay can be explained by shifts in the characteristics of employees. All other movements in wage growth, for example due to changes in productivity or inflation expectations, are considered 'unexplained'.

The effect of characteristics that are highly correlated can be controlled for using such regression techniques, which identify the marginal contribution of a specific factor. For example, in identifying the effect of industry on wages all other factors will be held constant, so managers with degrees and the same level of tenure are compared with similar individuals across industries, isolating the effect of working in different industries. This allows more precise estimates of the effect of changes in the mix of employment on pay.

How big are compositional effects typically and which factors have been driving them?

Changes in the composition of employment have typically boosted pay over time. On average these compositional effects added just under a ½ percentage point to annual wage growth between 1995 and 2015 (**summary chart**), which averaged around 3¼% over this period. This positive compositional effect is consistent with both the long-term trends in the workforce described previously, and other studies.(4)

The five factors which were found to make the largest contribution to these compositional effects were education, job tenure, age, occupation (the type of job the employee is doing) and industry. Other factors which were included in the analysis, but had much smaller aggregate effects, were gender, region, whether an individual worked full or part-time, whether they were employed in the public sector and whether they had a temporary employment contract.

The largest driver of this positive compositional effect over the past 20 years is the increasing education level of the UK workforce, which explains over 65% of the positive impact of compositional effects on wage growth. The other key contributor was a shift to more highly skilled occupations, which explains close to 30% of the positive effect. The

(1) These groups overlap so one individual may be counted in several averages shown.
(2) Quarterly LFS data is available from 1992, but these data sets do not have all the variables required for the analysis carried out here.
(3) Details of this process are provided in the annex.
(4) See, for example, Bell, Burriel-Llombart and Jones (2005).

Chart 6 Compositional effects on annual hourly wage growth relative to 1995–2010 average

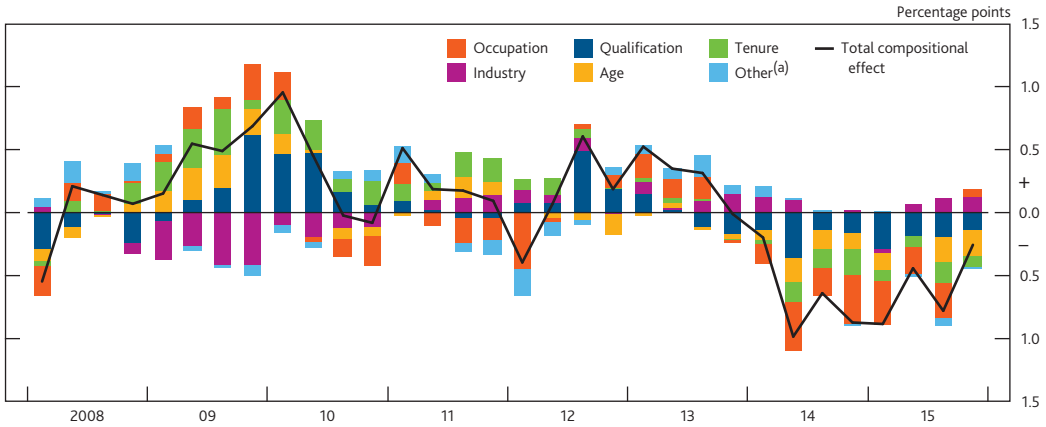

Sources: Labour Force Survey and Bank calculations.

(a) 'Other' includes gender, region, sector, working pattern and contract type.

changing industrial composition of the UK workforce is estimated to have acted as a net drag on wage growth.(1)

Compositional effects since the financial crisis

To understand the role of cyclical fluctuations in compositional effects, the estimates in a given period should be considered relative to the typical boost they provide. **Chart 6** shows the main results of the analysis, measuring compositional effects relative to their 1995–2010 average.(2) Bars above the zero line show a positive effect on wage growth for that factor relative to trend, with bars below the line indicating a negative effect. The 'total compositional effect' shows the overall effect of the different factors. By measuring compositional effects in this way, it is possible to determine whether they were having an abnormal effect on wage growth throughout the crisis and its aftermath — in particular in the past two years. At the start of the crisis, compositional effects pushed up strongly on wage growth by around 1 percentage point at their peak. In 2014 this effect went into reverse and compositional effects dragged down on wage growth by again roughly 1 percentage point. This may go some way to explaining the weak wage growth since that time shown in **Chart 1**.

The positive compositional effect at the beginning of the crisis was caused by a sharp decline in flows into employment from unemployment and a decrease in people retiring from the workforce. Those in unemployment typically have less education and are more likely to have worked primarily in lower-skilled industries. As such, the decline in the flow of the unemployed into employment meant that the average tenure, level of education and share of skilled occupations all increased substantially. In addition, some individuals are likely to have delayed retirement during the recession. As those close to retirement are likely to be more experienced and therefore be receiving higher pay, this too likely pushed up on the average wage level. These positive effects outweighed a drag caused by a decline in the number of those employed in relatively highly paid financial and insurance service industries jobs — shown by the purple bars in 2009–10.

As flows into employment recovered, many of these effects unwound, causing the compositional effect on wages to turn negative. For example, the average age of those who exited the workforce increased in 2014. It is possible that some of those who had previously delayed retirement due to the effects of the crisis were now exiting the labour force.

While these results are not conclusive, they are consistent with findings from other data. The Office for National Statistics (ONS) has found strong negative effects on mean wages in 2014 due to the difference between those leaving and entering employment.(3) Similarly, in other analysis(4) the ONS found a substantial increase in the 'quality' of the labour force at the beginning of the crisis and a sharp decline in 2014, where 'quality' seeks to capture factors such as qualification and age — in line with the results presented here.

In the latest data the impact of compositional effects on wage growth relative to average has fallen to around -¼% in 2015 Q4. Compositional effects on wage growth are likely to return to their long-run average as the labour market normalises.

Summary of the results

Table A summarises these results, showing the absolute effects different compositional factors have had on wage growth since the crisis. This shows that changes in

(1) It is important to note that the methodology used to calculate these compositional effects has limitations — most importantly it is limited by the available data and thus potentially exposed to bias by omitting other important variables. For example, the negative effect of industrial composition on wage growth could be due to certain industries being correlated with other factors such as union membership, which are not controlled for in the regression due to lack of quarterly data. If unions are able to increase worker pay, then a move to less unionised industries would show up in the result as industrial composition having a negative effect.
(2) These dates are chosen to try to include a full business cycle, while giving a stable benchmark against which to measure the compositional effects.
(3) See ONS (2015).
(4) See Connors and Franklin (2015).

qualifications and tenure boosted wage growth at the start of the financial crisis (second column) relative to the average (first column). The table also shows the drag compositional effects have had in 2014–15, driven by below-average contributions from qualifications, job tenure and occupation.

Table A Absolute contributions to wage growth from compositional changes

Per cent

	Average			
	1995–2010	2008–10	2011–13	2014–15
Total compositional effect	0.5	0.8	0.7	-0.1
Qualification	0.4	0.5	0.4	0.2
Tenure	0.0	0.2	0.1	-0.1
Age	0.0	0.1	0.0	-0.1
Industry	-0.1	-0.2	0.0	0.0
Occupation	0.2	0.2	0.2	0.0
Other	0.0	0.0	0.0	0.0

Compositional effects appear to explain a degree of the weak wage growth seen in 2014 and 2015; however they are not a complete explanation. Compositional effects explain little of the weakness of wage growth seen at the end of 2015 because the drag from compositional effects started to wane.

Given that compositional effects are also likely to affect productivity growth, some researchers have looked at compositional effects to see if they can help to explain the persistent weakness of productivity growth in the United Kingdom since the financial crisis, sometimes called the 'productivity puzzle'.(1) A box on page 19 looks at this relationship. It finds that compositional effects make the weakness in productivity early in the crisis more puzzling, but may help to explain some of the weakness since 2014.

Changing pay premium for different individual or job characteristics

A useful by-product of measuring compositional effects using the approach described above is that variations over time in the effect of different factors on wages can be estimated.

Of the five factors found to have the greatest effect on wages in the analysis — education, tenure, age, industry and type of occupation — the impact of three of these — education, tenure and age — have changed substantially over the 20-year sample period used for estimating compositional effects. **Chart 7** shows the effect of an individual's highest level of education on pay relative to having no qualifications at all, holding other factors constant. Since 1995, the effect of having a degree on pay has fallen substantially. In 1995, a degree would on average increase wages by 45% relative to having no qualifications at all; by 2015 this premium had fallen to 34%. Over the same period, the wage premium for A-levels and GCSEs also fell, but by far less.

Chart 7 The declining average effect of having a degree on wages

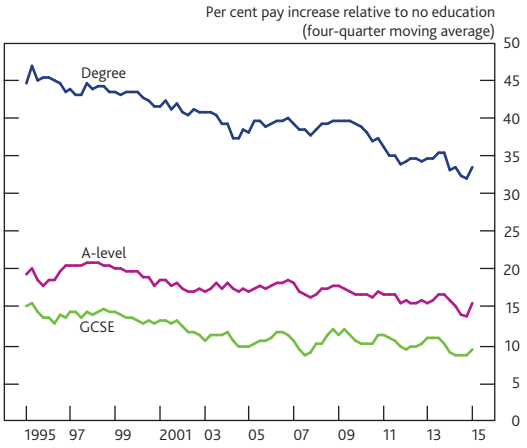

Sources: Labour Force Survey and Bank calculations.

There are a number of possible explanations which would be consistent with this finding. For example, if demand for highly skilled workers has not kept pace with an increase in supply, an increasing number of graduates would also lead to a decrease in the wage premium for those with degrees. Alternatively it is possible that the large increase in individuals studying for a degree in the United Kingdom has led to a fall in its signalling value (the ability of degrees to correctly identify more talented individuals) and thus the amount of pay which those with degrees can command.

The impact of job tenure on pay has also been declining over the past 20 years (**Chart 8**). While staying with an employer for a prolonged period of time does still have substantial benefits to an individual's earning potential, the premium for increased job tenure has declined. The premium for those with over 20 years' tenure has declined by around a third since 1995. This decline has happened without any significant shift in the tenure structure of the economy.

A third finding is that the relative earnings of those over 55 have increased since the financial crisis. Traditional age-earnings profiles have an 'inverted u' shape, earnings increase steeply with age up until the age of about 40, they are then stable until declining near the end of one's career. This is the profile found pre-crisis in the United Kingdom (as shown in the magenta bars in **Chart 9**). However, from 2008 onwards, this decline in the earning power of the United Kingdom's oldest workers was rapidly eroded and the results for 2015 show that there was no significant decline for any age group's earnings after age 35 (the orange bars in **Chart 9**). These statistical estimates cannot isolate the cause of this sudden change, and at this stage it is unclear how much, if any, of the change will prove to persist.

(1) See, for example, Blundell, Crawford and Jin (2013).

The productivity puzzle and compositional effects

This box provides a brief background on the post-crisis slowdown in productivity growth that has been seen in the United Kingdom, commonly known as the productivity puzzle. It also discusses to what extent compositional effects may help to explain this puzzle.

Since the financial crisis, labour productivity in the United Kingdom has been exceptionally weak. From 1997 to 2007 labour productivity (output per hour worked) grew at an average of 2.2% per year, whereas from 2008 to 2015 it grew at just 0.2% on average. This can be seen in **Chart A**.

Chart A Labour productivity growth declined after the crisis

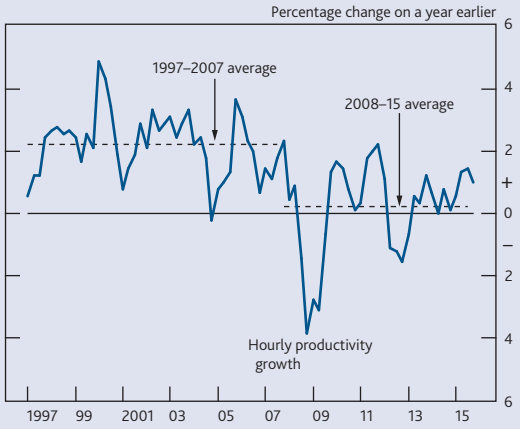

Sources: ONS and Bank calculations.

Measures of productivity can be used to inform estimates of an economy's ability to grow without generating excessive inflationary pressure. If the causes of this slowdown in productivity were primarily short term and cyclical, then the United Kingdom would likely be able to return to trend growth quickly as the economy recovered without it resulting in excessive inflation. On the other hand, a longer-term slowdown in labour productivity would point to a slower recovery, where wage growth, which is directly related to productivity growth, would be lower.

Most assessments of the productivity puzzle appeal to a number of factors to explain the apparent weakness.(1) These include some cyclical factors weakening productivity growth at the start of the crisis, for example as firms held on to workers at the cost of productivity.

But longer-term factors are also likely to have been important — especially those surrounding the impact of the financial crisis. Lower investment in capital, barriers to the efficient allocation of capital and the survival of lower-productivity firms throughout the crisis are all likely to have played some role in reducing productivity growth for longer than cyclical variation in productivity alone would suggest.

Given the likely direct impact of compositional effects on productivity growth it is worth examining whether they can play a role in explaining this puzzle. They do not provide an answer to the question of why productivity growth was substantially lower in the aftermath of the crisis. In fact compositional effects were likely to be pushing up on productivity growth between 2008 and 2013, making the weakness in productivity growth all the more puzzling.

More recently, however, compositional effects do appear to explain part of the below-trend productivity growth. **Chart B** shows productivity growth adjusted for compositional effects. With these adjustments recent productivity growth is closer to its pre-crisis trend. As the drag from compositional effects subside that could boost productivity growth.

Chart B Adjusting labour productivity for compositional effects

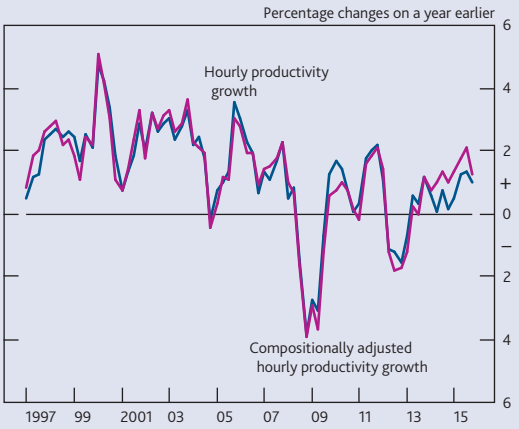

Sources: ONS and Bank calculations.

A reduced drag from compositional effects on productivity growth is one of the reasons the February 2016 *Inflation Report* projects increasing productivity growth in the medium term.

(1) See, for example, Barnett *et al* (2014).

Chart 8 The declining impact of tenure on wages

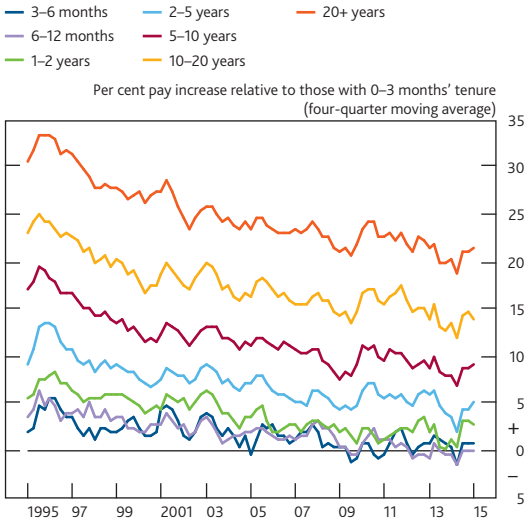

Sources: Labour Force Survey and Bank calculations.

Chart 9 The age-earnings profile relative to those aged 35–39

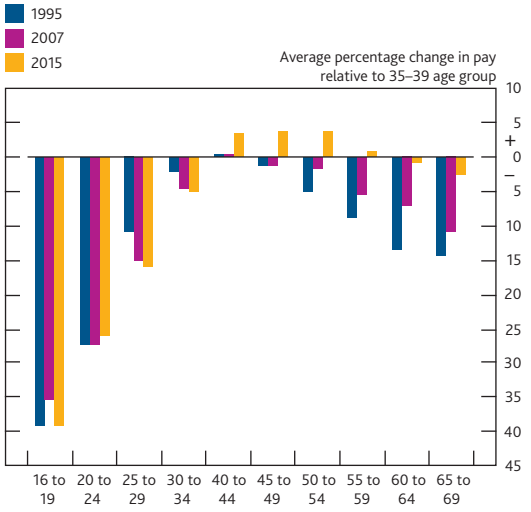

Sources: Labour Force Survey and Bank calculations.

While the effects of education, tenure and age have changed substantially, the role of occupation and industry have been much more constant over time. There is a large and stable premium in the earnings of high-skilled jobs relative to all others of between about 35%–50%. This contrasts with a much smaller gap between medium and low-skilled occupations of around 6% of earnings.

The pay gaps between industries are similar in size to those between occupations. Relative to a benchmark of manufacturing employees, sectors paying a premium include finance (19%), information, communications and technology (9%) and those in professional services (3%). In contrast, those in sectors such as health (-13%), education (-17%) or accommodation and food services (-21%) on average earn substantially less. In each case this is comparing workers with similar qualifications and roles across different industries. These gaps have been fairly stable over time.

While other factors did not play a significant role in the compositional effects on wages in the past 20 years, they still provide useful information. The pay gap between women and men, for example, has narrowed significantly over time according to these estimates. In 1995 women were paid 16% less than men for a given set of characteristics; by 2015 this figure had fallen to 9%. This narrowing of the gap likely explains why the increasing participation of women in the workforce has not been strongly pulling down on aggregate wage growth.

Wage differentials between different regions also appear, to a degree, to have converged. Relative to the South East, employees in London were paid on average 9% more in 1995; this has since declined to 7%. Meanwhile parts of the United Kingdom that experienced lower pay compared with the South East, such as the North of England, the West Midlands, Scotland and Wales, have all seen this pay gap narrow slightly.

Finally, these estimates suggest that the pay premium for those working in the public sector fell substantially in the late 1990s from a peak of 10% relative to those in the private sector in 1996, to around 4% in 2002. Since then it has remained relatively unchanged.(1)

Conclusion

Changes in characteristics such as education, occupation and age can have a significant effect on wage growth. Over the past 20 years, as the United Kingdom's workforce has become more educated, moved to higher-skilled roles and aged, these effects are estimated to have pushed up on annual wage growth by an average of roughly ½ a percentage point.

Following the financial crisis, due to large changes in the flows into and out of employment, these compositional effects became particularly pronounced. Initially, they pushed up on wage growth by up to a percentage point relative to average as high earners stayed in employment longer than expected and lower earners became unemployed. At the end of 2013 these shifts went into reverse and compositional effects pushed down wage growth. This helps to explain some of the weakness in wage growth that was seen in 2014 and 2015.

More recently, data for 2015 Q4 show that compositional effects have started to subside. The drag on wage growth is likely to dissipate as the labour market normalises and the effect on wages of changes in the composition of the workforce returns to normal.

(1) The ONS notes that public sector workers typically work for larger employers which provide a pay premium relative to smaller employers. Once this is taken into account, the ONS finds that public sector workers earn less than their private sector counterparts. This highlights the sensitivity of some of our results to the controls that are used. See ONS (2014).

Annex
Calculating compositional effects

In order to calculate the compositional effect on wage growth between two given periods a Oaxaca decomposition is used, following the approach used by Blundell, Crawford and Jin (2013).

This decomposition splits the wage change between two time periods, between observed components, such as work experience, tenure and industry and a residual that cannot be explained by these observable characteristics. In this analysis, this unobserved component is a proxy for factors such as technological change, which are likely to push up on productivity and pay of the entire workforce over time. Changes in wages between two time periods, which can be explained by changes in observed characteristics, are recorded as 'compositional effects'.

Formally, the estimation is based on a linear model where wages at time t are determined by a vector of characteristics X_t plus an error term:

$$Y_t = X_t' B_t + e_t$$

The change in wages between time t and $t+n$ can be written as:

$$E(\Delta Y_{t+n}) = \textit{Compositional effect} + \textit{Coefficient effect} + \textit{Interaction term}$$

Where these specific effects can be written as:

$$\textit{Compositional effect} = (E(X_{t+n}) - E(X_t))' B_t$$

$$\textit{Coefficient effect} = E(X_{t+n})'(B_{t+n} - B_t)$$

$$\textit{Interaction term} = (E(X_{t+n}) - E(X_t))'(B_{t+n} - B_t)$$

The coefficient effect and interaction term are combined and treated as an unexplained component of wage growth in the analysis presented.

The specific linear model which is estimated is the following:

$$\begin{aligned}\ln(\textit{hourly\_wages})_i &= \textit{gender}_i + \textit{highest\_qualification}_i \\ &+ \textit{industry}_i + \textit{age\_group}_i + \textit{tenure}_i \\ &+ \textit{region}_i + \textit{occupation}_i \\ &+ \textit{public\_sector}_i + \textit{full\_time}_i \\ &+ \textit{temporary\_contract}_i\end{aligned}$$

All of the variables in the regression, with the exception of wages, are discrete variables which are programmed as a series of dummy variables described in **Table A1**.

Table A1 Characteristics included in compositional effects regression

Variable	Categories
Qualification	Degree; Higher education; A-level or equivalent; GCSE or equivalent; Other qualification; No qualification; Don't know.
Industry	SIC2007 Industry section codes.
Age groups	16–19; 20–24; 25–29; 30–34; 35–39; 40–44; 45–49; 50–54; 55–59; 60–64; 65–69; 70+.
Tenure	< 3 months; 3–6 months; 6–12 months; 1–2 years; 2–5 years; 5–10 years; 10–20 years; 20+ years.
Region	North; Yorkshire & Humber; East Midlands; East Anglia; London; South East; South West; West Midlands; North West; Wales; Scotland; Northern Ireland.
Occupation	Managers, directors and senior officials; Professional occupations; Associate professional and technical occupations; Administrative and secretarial; Skilled trades; Personal services; Sales and customer services; Process, plant and machine operatives; Elementary occupations.
Gender	Male; Female.
Working pattern	Full-time; Part-time.
Sector	Private sector; Public sector.
Contract type	Permanent; Temporary.

Hourly wages are used to abstract from the issue of varying average hours worked in the economy.

The model is estimated on LFS quarterly data from 1994 Q1–2015 Q4 where individual observations are dropped if any individual characteristics used in the regression are missing.

References

Barnett, A, Batten, S, Chiu, A, Franklin, J and Sebastiá-Barriel, M (2014), 'The UK productivity puzzle', *Bank of England Quarterly Bulletin*, Vol. 54, No. 2, pages 114–28, available at www.bankofengland.co.uk/publications/Documents/quarterlybulletin/2014/qb14q201.pdf.

Bell, V, Burriel-Llombart, P and Jones, J (2005), 'A quality-adjusted labour input series for the United Kingdom (1975–2002)', *Bank of England Working Paper No. 280*, available at www.bankofengland.co.uk/research/Documents/workingpapers/2005/wp280.pdf.

Blundell, R, Crawford, C and Jin, W (2013), 'What can wages and employment tell us about the UK's productivity puzzle?', *IFS Working Paper W13/11*, available at www.ifs.org.uk/wps/wp201311.pdf.

Connors, E and Franklin, M (2015), 'Quality adjusted labour input, estimates to 2014', available at www.ons.gov.uk/ons/dcp171766_402986.pdf.

Leaker, D (2008), 'The gender pay gap in the UK', *ONS Economic and Labour Market Review*, Vol. 2, No. 4, pages 19–24.

Office for National Statistics (2014), 'Public and private sector earnings — November 2014', available at www.ons.gov.uk/ons/dcp171776_383355.pdf.

Office for National Statistics (2015), 'Understanding average earnings for the continuously employed — using the annual survey of hours and earnings 2014', available at www.ons.gov.uk/ons/dcp171766_404125.pdf.